Theoretical Knowledge and Inductive Inference

John-Michael Kuczynski

Chapter 1 Inferential knowledge and logical dependence
Two kinds of knowledge: direct and indirect
Two kinds of dependence-relations: causal and logical
Propositions the objects of belief
Relations of logical dependence not known through observation
Analytic truth and analytic knowledge
Analytic knowledge is always knowledge of dependence-relations.
Knowledge vs. awareness
Truths, not objects, the objects of explanation
Causal relations known through observation
All theoretical knowledge inferential but not vice versa
Theories and theoretical knowledge
All knowledge justified true belief but not vice versa
Propositions the objects of belief
Knowledge vs. awareness
Causal vs. logical necessity
The concept of evidence
Causal series, structure-preserving series, and information-bearing series
Information≠
evidence
Information-transmission as the basis of knowledge
Causation the essence of temporal precedence
Continuity the essence of causation and information-transmission

Chapter 2 Deduction, induction, and the conditional nature of scientific knowledge
Entailment and deductive inference
Confirmation and inductive inference
Four principles of logic
Entailment≠
maximal confirmation
Confirmation-relations non-objectual in nature
Conditionals
Laws of nature often given by universal generalizations
The scope of natural law not confined to actual or even possible objects
Knowledge of the imperceptible identical with knowledge of conditional truth
Three kinds of imperceptibles
The epistemic function of natural law
Laws vs. singular truths
Laws sometimes given by place- and time-sensitive truths
Meaningful generalizations have a systemic basis

Chapter 3 Causation
Causation a relation between events, not objects
Conditions identical with event-patterns
So-called changes typically identical with meta-changes
Explanation, prediction, and control
Stable-event patterns a prerequisite for explanation-identification
Freud on causation in relation to the development of mental illness
Psychological overdetermination
Threshold effects and the analog-digital distinction
Counterfactual conditionals
David Lewis's analysis of counterfactuals
Program-causality
Causal redundancy in relation to counterfactuals
The Counterfactual Analysis of Causality (CAC)
The motivation for CAC
CAC evaluated
Problems with CACD
(i) Explained

(ii) explained
(iii) Explained
Summary of our discussion of The Counterfactual Analysis of Causality (CAC)

Chapter 4 The Regularity Analysis of Causation
The Regularity Analysis of Causation stated
Hume's argument for this analysis (first version)
Hume's argument for this view (second version)
The fallacy embodied in Hume's analysis—Part 1
The fallacy embodied in Hume's analysis—Part 2
The fallacy embodied in Hume's analysis—Part 3: The Concept of a Hume-world and its uselessness in the way of supporting Hume's analysis
Hume's argument for this claim
Hume worlds observationally different from their non-Humean counterparts
Three additional objections to the Regularity Analysis of Causation
Explanation of Problem #1
Explanation of Problem #2
Explanation of Problem #3

Chapter 5 Causation to be understood in terms of the shortcomings of the deductive-nomological model of scientific explanation
The deductive-nomological model of explanation (DN)
Problems with DN
Nomic vs. accidental generalizations
How DN-advocates deal with the distinction between nomic and accidental generalizations
DN viciously circular if it requires laws to support counterfactuals
Additional DN-internal vicious circularity
Some problems with this analysis
Causal knowledge as differential knowledge
Causation as proportionality restoration
Physical laws as regularities of proportion
Explanation of the fact that statistics tend to track causal substructure
The instantaneousness of the interval between cause and effect
Natural Laws dispositive of, but not identical with, regularities
How we know what causes what

Knowledge-why=high-resolution knowledge-what
Singular causation and the falsity of the Deductive-Nomological Model of explanation
The failure of the Deductive-Nomological model in connection with psychological explanation
DN in relation to statistical explanation
Delimiting DN's sphere of application
The Raven-Paradox
Hempel's solution
The problem with Hempel's solution
The actual solution to the Raven Paradox
The Raven-paradox a by-product of the Regularity Analysis

Chapter 6 The Problem of Induction
What is the problem of induction?
The problem with the problem of induction
No such thing as enumerative induction
(1) Defended
(2) Defended
Discontinuity-elimination the essence of explanation
Inductive knowledge as knowledge of continuities of continuities
The justification for Inference to the Best Explanation
The justification for Inference to the Best Explanation (continued)
Finalizing the justification for Inference to the Best Explanation
The Concept of Natural Law in relation to the problem of induction
Objects as causal series

Chapter 7 Determinism, implementation-indeterminism, and the nature of natural laws
"Determinism" defined
The first reason why determinism doesn't entail complete predictability: Interference effects
Laplace's Thesis
Why the existence of the unknown must be granted
Second reason why determinism doesn't entail complete predictability: No way to predict how knowledge will grow
Supervenience
Supervenience≠identity
Psychological supervenience difference from other kinds of supervenience

Bound by physical law≠
explained by physical law
Software-determinism≠
statistical determinism
Implementation-determinism≠
program-determinism
Probability in relation to higher-order mechanisms
Statistical trends reflective of, but distinct from, probabilities
Evidence of probabilities vs. actual probabilities
Conclusion of Chapter 7

Chapter 8 Measurement
Measurement vs. enumeration
Relativization to a standard the essence of measurement
Relative vs. absolute length
Three axioms of measurement-theory
Scalar quantities vs. vector quantities
Scalar magnitudes identical with one-dimensional manifolds
Congruence
n-dimensional congruence
The use of provisional metrical standards and approximate physical laws
Triangulation and precisification
Precisification in relation to magnitudes other than length
Summary of findings
Synchronic vs. diachronic measurements
Metrical Conventionalism
Extensive vs. intensive magnitudes
Relativity Theory and the non-additivity of scalar quantities
Direct measurements always yield rational numbers
Nomic Conventionalism
This argument refuted
Conventionalism vs. Strict Conventionalism
Conventionalism inconsistent with the desideratum that causal anomalies be eliminated
Conclusion of the present work

Chapter 1 Inferential knowledge and logical dependence

Two kinds of knowledge: direct and indirect

Some knowledge is direct; some is indirect.
Indirect knowledge is knowledge that is acquired through inference.
To make an inference is to form a new belief on the basis of an old one.
Inferences, when legitimate, are made in accordance with legitimate rules of inference.
Rules of inference, when legitimate, correspond to dependence-relations.

Two kinds of dependence-relations: causal and logical

There are two kinds of dependence-relations: logical and causal.

Logical dependence-relations hold among *propositions*. A proposition is a truth or a falsehood.

Example of a logical dependence-relation: The proposition that x is a triangle cannot be true unless the proposition that x has three sides is true. Therefore, x's being a triangle depends on x's having three sides.

Causal relations hold among states of affairs, not truths and falsehoods. (A state of affairs is anything that is in space-time.)

Example of a causal dependence-relation: Smith cannot pass the class unless Smith studies. Therefore, Smith's passing depends on his studying.

The first-dependence relation is interpropositional, i.e. it holds between propositions.

The second is objectual, i.e. it holds between occupants of the space-time manifold and therefore between 'objects,' relative to some delineation of that term.

Propositions the objects of belief

The objects of knowledge are propositions. It isn't known unless it's believed and it isn't believed unless it's a proposition.

One can be *aware* of non-propositions, e.g. rocks, trees, people, but one does not in the relevant sense *know* them. Rather, one knows truths *about* them.

Relations of logical dependence not known through observation

Some dependence-relations are known, not through sensory observation, but through conceptual analysis. Such truths are *analytic*, as opposed to empirical. Examples of analytic truths are:

1. If P, then either P or Q.

2. There are infinitely many primes.

3. Nothing can be numerate or literate without being sentient.

4. There are no laws where there is no government.

5. There is no recursive definition of the class of Dedekind-cuts.

6. 1+1=2.

7. 1+1≠3.

1 obviously expresses a dependence-relation, and 2-7 non-obviously do so, as their meanings are:

2. A set's having finitely many members depends on its not having the set of prime numbers as a subset.*

3. An entity's being literate or numerate depends on its being sentient.*

4. The presence of law depends on the presence of government.*

5. A set's having either finitely many or denumerably many members depends on its not having a cardinality equal to or greater than the cardinality of the class of Dedekind cuts.*

6. A set's being a couple depends on its being the union of two non-overlapping unit sets.*

7. A set's being a triple depends on its not being the union of two non-overlapping unit sets.*

Analytic truth and analytic knowledge

Analytic knowledge is always knowledge of dependence-relations.
Empirical knowledge can be either direct or indirect; the same is true of analytic knowledge. [1]

Knowledge vs. awareness

Knowledge is conceptually articulated awareness. In order for me to know that my shoes are uncomfortably tight, I need to have the concepts *shoe*, *tight*, *discomfort*, etc. I do not need to have these concepts---or, arguably, *any* concepts---to be *aware* of the uncomfortable tightness in my shoes. My knowledge of that truth is a conceptualization of my awareness of that state of affairs.

Equivalently, there are two kinds of awareness: propositional and objectual. My visual perception of the dog in front of me is a case of objectual awareness, as is my awareness of the tightness of my shoes. My knowledge that there is a dog in front of me is a case of proposition-awareness, as is my knowledge that my shoes are uncomfortably tight.

Truths, not objects, the objects of explanation

Observations are objectual awarenesses. The contents of such awarenesses must be converted into propositions if they are to be explained. This is because it is truths that are explained, and truths are true propositions.

"But don't we explain events?" it will be objected. "Don't we explain bolts of lightning and avalanches?"

To explain some avalanche or lightning-bolt x is to explain why it is a truth *that x occurred.* One is aware of lightning-bolts and avalanches. One explains the corresponding truths.

Causal relations known through observation

Causal relations can be known only through one's senses (through sight, touch, etc.) and through real-time awareness of one's own psychological processes. In other words, causal relations can only be known *empirically*; and knowledge of them is therefore *empirical.*

All theoretical knowledge inferential but not vice versa

Theoretical truths are necessarily known through inference, and knowledge of such truths is therefore inferential.

But not all inferential knowledge is theoretical knowledge. I notice that you look tired and on that basis know that you didn't get enough sleep. That knowledge is inferential but non-theoretical.

Theories and theoretical knowledge

A theory is given by a proposition that, if true, describes some relatively comprehensive fact about the structure of the universe.

Theoretical knowledge is either

(i) Knowledge of such a theory (e.g. your knowledge of evolutionary theory)

or

(ii) Knowledge that is derived from knowledge of such a theory (e.g. your knowledge that, the laws of mechanics being what they are and projectile x's state of motion being what it is, x will collide with surface y in approximately two hours).

All knowledge justified true belief but not vice versa

If you know it, then

(i) You believe it,
(ii) It is true, and
(iii) Your reasons for believing it are legitimate.

At the same time, given only that it's true and that you are justified in believing it, it doesn't follow that you know it. A broken clock is right twice a day. If your belief that it's noon is based on the read-out of a broken clock, it isn't knowledge, even if that clock happens to be right and, in addition, you know it to be relatively reliable.[2]

Your belief *would* have been knowledge had you acquired it by virtue of being *en rapport* with the relevant realities; and the reason it wasn't knowledge is that it was not so acquired.

For you to be *en rapport* with those realities is for you to be linked to them by structures that are *information*-transmissive, as opposed to merely *evidence-generative*. The italicized terms will be presently defined and the justification for the just-made assertion will thereby become clear.

Propositions the objects of belief

The objects of knowledge are propositions. (It isn't known unless it's believed and it isn't believed unless it's a proposition.) One can be *aware* of non-propositions, e.g. rocks, trees, people, but one does not in the relevant sense *know* them. Rather, one knows truths *about* them.

Knowledge vs. awareness

Knowledge is conceptually articulated awareness. In order for me to know that my shoes are uncomfortably tight, I need to have the concepts *shoe*, *tight*, *discomfort*, etc. I do not need to have these concepts---or, arguably, *any* concepts---to be *aware* of the uncomfortable tightness in my shoes. My knowledge of that truth is a conceptualization of my awareness of that state of affairs.

To put it another way: There are two kinds of awareness: propositional and objectual. My visual perception of the dog in front of me is a case of objectual awareness, as is my awareness of the tightness of my shoes. My knowledge that there is a dog in front of me is a case of proposition-awareness, as is my knowledge that my shoes are uncomfortably tight.

Causal vs. logical necessity

Causal necessity is different from logical necessity. Given that x is a triangle, it is *logically* necessary that x have three sides. Given that x is an ice-cube that was just placed on a very hot surface, it is *causally* necessary that x begin to melt.

The concept of evidence

One state of affairs e is evidence of another state of affairs e* iff (if and only if) e and e* are causally related in such a way that, if one knows of e's existence, one has reason to believe in e* existence. If one knows that Smith is coughing up green phlegm, one has reason to believe that Smith has a lung infection. Therefore, Smith's coughing up phlegm is evidence of Smith's having a lung infection.

Causal series, structure-preserving series, and information-bearing series

S is a *causal series* if S consists of events $e_1 \ldots e_n$ such that, for any i ($1 \leq i \leq n$), e_i is the cause of e_{i+1};

A *structure-preserving* series if e_1 transmits its structure to e_n;

An *information-bearing* series if a structure-preserving series.

Information-transmission is structure-preservation. When you talk on the phone, the structure of the sounds you make is coded into electrical signals. The other person's phone converts those signals into isomorphic (structure-identical) sounds. The second sound-class therefore bears information about the first.

Information≠ evidence

All information is evidence, but not all evidence is information.

The elevator's stopping at the 5^{th} floor is evidence of somebody's pushing a certain elevator-related button. But neither event bears information about the other, owing to the fact that, since they have different structures, the first didn't transmit its structure to the second.

Information-transmission as the basis of knowledge

Our sense-perceptions are isomorphs of the objects they represent: the latter transmit their structures to the former.

A perception-like experience is veridical, and therefore a *bona fide* perception, just in case it inherits its structure from its object, and accurate *to the extent* that it does so.

It is only to the extent that our sense-perceptions, taken both individually and collectively, inherit the structures of their external counterparts that our perception-based beliefs are accurate.

It is only to the extent that our beliefs are linked by information-transmissive series to the corresponding realities that those beliefs are knowledge.

A testimony-based belief is knowledge only if the testimony in question is constitutive of an information-transmissive link to the externality in question.

A belief based on the testimony of a dishonest person is knowledge, so long as the testimony in question was information-transmissive. A belief based on the testimony of an honest person is not knowledge unless the testimony in question was information-transmissive.

A belief based on an unreliable mechanism (e.g. a clock that often stops) is knowledge, so long as, in the relevant context, the mechanism was information-transmissive. A belief based on a reliable mechanism (e.g. a clock that seldom stops) is not knowledge unless, in the context in question, the mechanism was information-transmissive.

The difference between a clock that seldom stops and one that often stops is purely statistical. The same is true of the difference between a witness who seldom lies and one who often lies.

The difference between knowledge and mere belief is not statistical. A given belief is knowledge exactly if it is the end-result of an information-transmissive series, it being irrelevant how other beliefs originate.

Causation the essence of temporal precedence

For one event e to precede another event e* is for there to be a possible causal process beginning with e and ending with e*.

For e and e* to be simultaneous is for there to be no such process.

Continuity the essence of causation and information-transmission

Effects are continuations of their causes.

An event transmits information about its cause to the extent that the former inherits its structure from the latter and, therefore, to the extent that the latter's structure is causally determinative of the former's.

When the cause-effect relation preserves structure—that is, when information is transmitted---it is because multiple concurrent causal lines are preserving their mutual relations. When a high-resolution of someone's face is transmitted, without loss of detail, from one computer to another, each of the facial features represented generates its own causal series *and* the mutual relations of those various causal series are preserved. Consequently, circumstance-specific knowledge of causal dependencies isn't needed to infer the nature of the cause from the nature of the effect. But just such *ad hoc* knowledge *is* needed to infer that somebody pushed the doorbell-button from the ringing of the button.

Knowledge is initially acquired through information-transmission series---our thoughts inherit their structures from our perceptions, which inherit their structures from the corresponding externalities. On the basis of this knowledge, we can acquire knowledge of *non*-information-transmissive evidential relations. But knowledge of the latter kind presupposes knowledge of the former kind.

Chapter 2 Deduction, induction, and the conditional nature of scientific knowledge

Entailment and deductive inference

Given two propositions P and Q, P *entails* Q iff Q is a logical consequence of P.

If P entails Q, then, if P is true, it follows that Q is true and therefore that Q may be inferred from P.

If a given person infers Q from P, and he does so on the basis of his knowledge that P entails Q, then that has person has made a legitimate *deductive* inference.

Confirmation and inductive inference

P *confirms* Q iff P, supposing it true, increases the likelihood that Q is true.
Given only that P confirms Q, it doesn't follow that Q may be inferred form P.
There are three reasons for this.
First, if P confirms Q, P may also confirm not-Q.

P: Smith is an obese chain-smoker.
Q: Smith dies of a coronary, not of cancer.
R: Smith dies of cancer, not of a heart-attack.

Q is incompatible with R. P confirms each, and we may suppose for argumentative purposes that it confirms each equally. Therefore, given P, it doesn't follow that Q may be inferred and it doesn't follow that R may be inferred. (Nor does it follow that *Q or R* may be inferred, given P may be confirm propositions incompatible with each other Q and R.)

Also, given only that P confirms Q, it doesn't follow that P, supposing it true, makes it probable that Q is true. What follows is merely that, because P is true, Q is more probable than it would otherwise be.

P: Smith is extremely intelligent.
Q: Smith will find a cure for cancer.

P confirms Q, since the probability of Q, given P, is higher than the probability of Q, given not-P. But the first probability is still very small. Therefore, P does not by itself warrant acceptance of Q.

Third, supposing that

(i) P is true;

(ii) P makes it highly probable that Q;

and

(iii) P doesn't equally confirm some Q-incompatible proposition R;

nonetheless,

(iv) Q *still* cannot be inferred from P.

This because there might be some true proposition R that, when coupled with P, disconfirms Q. Consider:

P: Smith is an extremely smart person who very much wants to become a lawyer and is about to get his degree from a top law school.
Q: Smith will be a successful attorney.
R: Smith's country is about to be taken over by a dictator who intends to outlaw the practice

of law.

Supposing that P confirms Q, Q may legitimately be inferred from P if, and *only* if,

(a) The supposition that P is true and Q is false has fewer unexplainable consequences than the supposition that P and Q are both true,

and

(b) The supposition that P is true has fewer unexplainable consequences than the supposition that P is false.

If a given person infers Q from P, and he does so on the basis of his knowledge of each of (a) and (b), then that person has made a legitimate *inductive* inference.

Four principles of logic

(i) Entailment is transitive: if P→Q, and Q→R, then P→R.

P: Smith has over 20 cars.
Q: Smith has over 15 cars.
R: Smith has over 10 cars.

(ii) Confirmation is not transitive: If P confirms Q, and Q confirms R, it doesn't follow that P confirms R. To wit,

P: Smith is a prisoner who is due to be executed in 20 minutes and whose medical condition is such that, if a competent physician who didn't know about Smith's upcoming were to examine Smith, that physician would give Smith at least another year to live.
Q: If a competent physician were to examine Smith, he would conclude that Smith would not die a natural death for at least nother year.
R: Smith will live at least another year.

(iii) Entailment is monotonic: If P→Q, then, for any R, (P and R)→Q.

P: Smith has over 20 cars.
Q: Smith is from Sudan.
R. Smith has more than 10 cars.

(iv) Confirmation is non-monotonic: Given only that P confirms Q and that R is true, it doesn't follow that P and R jointly confirm Q.

P: The match is struck.
Q: The match ignites.
R: The match is damp.

Entailment≠
maximal confirmation

If P entails Q, then, supposing that P is true, there is a 100% probability that Q is true.

But entailment is not maximal confirmation.

Entailment is a logical notion. Confirmation is an explanatory notion.

For P to entail Q is for *P and not-Q* to be counterlogical. For P to confirm Q is for *P and not-Q* to be counterexplanatory.

A given proposition is counterlogical if its very own structure prohibits it from being true.

A given proposition is counterexplanatory if, given the available data, the supposition that it's true has more unexplainable consequences than does the supposition that it's false.

Confirmation-relations non-objectual in nature

Confirmation relations, though *not themselves* causal relations, sometimes hold *in virtue of* causal relations.

P: Smith smokes 10 packs of cigarettes every day.
Q: Smith has high blood pressure.

P confirms Q, the reason being that Smith's smoking so much is likely to cause him to have high blood pressure. But neither P nor Q is a cause or an effect, and P's confirming Q therefore isn't an instance of causation.

Conditionals

A conditional is a statement of the form: *if P, then Q.* An example of a conditional is:

(1) If Smith has more than 100 houses, then Smith has more than 50 houses.

Sentences beginning with expressions like "supposing that," "assuming that," "since," "given that," "on the condition that," and "because" are also conditionals ("since/assuming that/given that/on the condition that Smith has more than 100 houses, he has more than 50 houses").

These terms are not all synonymous, but the semantic (meaning-related) differences between them are not important in this context.

Laws of nature often given by universal generalizations

(1) concerns a *specific* person. It is a *singular* proposition and therefore, if true, a *singular* truth. Laws of nature tend to be expressed by *general* truths; and they tend to be expressed by statements to the effect that each member of an indefinitely large *class* of conditionals is true.

Consider the law that:

(2) Copper expands when heated

(2) 's meaning is perspicuously given by the statement that:

(3) Given any object x, if x is made of copper, then x expands when heated.

(3) does not say of any *one* object that it expands when heated. (3) concerns an entire class of objects; for it says that, for *any* object A,

(4) If A is made of copper, then A expands when heated.

Thus, whereas (1) entitles you to speak only about Smith, (3) entitles you to speak about *anything*. Given any copper object, (1) entitles you to say that it will expand if heated. And given *any object it at all*, (1) entitles you to say of that object that

(a) It would expand if it were copper;

or supposing that it *is* in fact copper but that it isn't being heated, that

(b) It will expand if heated;

or supposing that it *is* being heated *and* is made of copper, that

(c) It is expanding because it is copper that is being heated.

The scope of natural law not confined to actual or even possible objects

Thus, laws of nature have a great deal of scope. They are not restricted to this or that object; they hold with respect to entire classes of objects; and, what is more, those classes sometimes comprise merely *possible* objects.

In fact, many laws of nature hold *only* with respect to classes all of whose members are *merely* possible (meaning that, in actuality, they don't have members but could have members). For example,

(5) An object undergoes no changes in shape or size if the forces acting on it cancel each other out.

Further, many such laws hold with respect to classes whose members are *prohibited by the laws of nature themselves* from existing. For example,

(6) An object undergoes no changes in shape or size if no forces are acting on it.

In fact, laws like (6) are especially important, since they describe limiting cases in terms of which actual cases are to be understood. Even though the laws of nature prohibit the existence of an objects on which *no* differential forces are acting, it is only *to the extent* that a given object's condition approximates to that of such an object that it is shape- and size-stable.

Knowledge of the imperceptible identical with knowledge of conditional truth

Knowledge of the future is conditional. Given (3), and given also that

(7) O is copper that is being heated but is not otherwise subject to differential forces,

it follows, and can therefore be known, that

(8) O will soon be bigger than it is now.

Knowledge of the past is conditional:

(9) Given that Jerry's ice-cream cone is melting on the table-top, said ice-cream cone was removed from a cold environment more than 10 seconds ago.

Knowledge of what is possible but non-actual is conditional:

(10) Jerry's (non-melting because freezer-internal) ice-cream cone would melt if the freezer were unplugged.

Finally, knowledge of what is present but imperceptible is conditional:

(11) Given that Jerry's (next-room-located and therefore imperceptible to me) ice-cream cone is melting, it stands to reason that it was recently removed from an n°≤32 environment.

Three kinds of imperceptibles

Among the various ways in which physical objects can be imperceptible, three are of particular importance.

First, they can be imperceptible to a given person, in a given context, owing to the fact that, in that context, the causal (information-bearing) processes that mediate between subjects and otherwise perceptible objects have been severed. (Cf. the ice-cream cone that is in the next room. Any light-rays bouncing off of it are intercepted by the opaque walls of that room before they can reach you.)

Second, they can be imperceptible because they are too small to see or otherwise perceive, even though, given the right technology, e.g. a microscope of the right kind, they *could* be perceived. (Cf. cells, viruses, and molecules.)

Third, there are some physical objects which, for reasons of natural law, could not possibly be perceived, no matter how much technology improved. (Cf. photons, subatomic particles.)

The epistemic function of natural law

Scientists aspire to identify natural laws on the basis of which they can, when given context-specific information, make true statements as to future, past, or possible behavior of that context as well as other contexts.

Laws vs. singular truths

Any interest that science has in a given fact is subordinate to its interest in the principles in terms of which that fact is to be understood.

The scientist doesn't care whether he is studying this as opposed to that penguin/flower/projectile, except in so far as the one penguin/flower/projectile instantiates general principles/natural laws not instantiated by the other.

Laws sometimes given by place- and time-sensitive truths

Some philosophers of science have taken the term 'general principle' to mean 'spatiotemporally invariant principle,' and they have thus taken the position that the laws of nature must be spatiotemporally invariant. A principle/law is 'spatiotemporally invariant' if, given any two regions of space-time, it holds no less for the one than the other.

But many of the dependence-relations on which scientists depend are undefined, and therefore neither true nor false, for almost all regions of space-time. Given that you are human, it can be inferred that you will fall ill if you consume cyanide, but not that you will go into anaphylactic shock if you eat shellfish. Consequently,

(a) x is human→x is cyanide-intolerant

is a legitimate rule of inference, whereas

(b) x is human→x is shellfish-intolerant

is not a legitimate rule of inference.

But even though the class of space-time regions in which there are people is a vanishingly small sub-class of the total class of space-time regions, and even though, consequently, the class of space-regions with respect to which (a) is non-vacuously true is correspondingly small, (a) is nonetheless *ultimately* in the same category as a paradigm-case of a law, such as

(c) nothing exceeds the speed of light.

Each of (a) and (c) supports predictions and counterfactuals; (b) does not do so, and neither does an accidental, albeit place- and time-invariant, generalization, such as

(d) nothing in any region of space-time that is made of solid gold weighs exactly n oz.,

for an appropriate value of n.

The difference between (a) and (c) is one of degree: (a) is more approximate than (c) and (a) also has less scope than (c). But with those qualifications, (a)'s role in the acquisition of knowledge is analogous to (c)'s role: each serves as a generalization on the basis of which specific situations can be evaluated. This cannot be said of (b) or of (d).

In order to discover the structure of the world, we must organize the data at our disposal; and in order to organize the data at our disposal, we need to use principles, such as (a), that are of limited scope and are only approximately true.

Suppose *arguendo* that, were there ever nothing left to know about the world, there would be no place for low-grade generalizations like (a). But were that condition met, there won't be any need for science; and so long as it isn't met, principled but defeasible generalizations, like (a), are to be co-categorized with principled, *non*-defeasible generalizations, like (c), i.e. they are to categorized as *laws of nature*.

Also, that condition cannot possibly be met. With every increase in our knowledge, there is a change in what there is to know about the world, given that any given instance of knowledge is no less a constituent of the world than is a rock or a chair. Therefore, attempts to plug up the

holes in our knowledge of the world create new such holes, it being the purpose of science, not to eliminate such holes, but to minimize their number and to restrict the scope of those cannot be eliminated.

Meaningful generalizations have a systemic basis

The difference between

(a) x is human→x is cyanide-intolerant

and

(b) x is human→x is shellfish-intolerant

is not to be understood in statistical terms. While it is presumably a fact that 100% of human beings are cyanide intolerant, whereas only a small percentage are shellfish intolerant, this isn't the reason that (a) is, whereas (b) is not, a law.

It is inherent in what it is to be physiologically human (or otherwise mammalian) that a human being be cyanide intolerant. The reason that no one is cyanide-tolerant is that a person's being cyanide-*in*tolerant is rooted in his having a distinctively mammalian physiological make-up.

And if a given person happened to be cyanide-tolerant, he would be so *despite* his being physiologically human. It follows that if *everyone* happened to be cyanide-tolerant, while otherwise being physiologically just like actual human beings, each person would be cyanide-tolerant *despite* his being physiologically human; and it therefore follows that, under such circumstances, each and every instance of cyanide-tolerance, though statistically normal, would in constitute a singularity.

By the same token, if everyone happened to be shellfish-intolerant, while otherwise being just like a physiologically normal human being, none of these instances of shellfish-intolerance would be rooted in what it is to have a distinctively human physiology.

At the same time, even though shellfish-tolerance isn't counterphysiological, neither is shellfish-intolerance, this being why a case of shellfish-intolerance in a world of otherwise shellfish tolerant human beings would represent a singularity only in some emptily statistical sense.

Laws of nature are not perfect concomitances: for it to be a law of nature that instances of phi bring about instances of psi is not for each and every phi to be followed by an instance of psi. There could be such a law even if *no* phi every brought about any psi. For there to be such a law is for instances of phi to be dispositive, for reasons of a relatively non-context-specific nature, of the occurrence of instances of psi.

These statements will be further defended in Chapter 5.

Chapter 3 Causation

Causation a relation between events, not objects

If x causes y, then x and y are events. They are not are not rocks, trees, or snowmen. They are not *things*, in other words, at least not in the sense in which we will use that word.

Though not themselves causes or effects, things have *causal properties.* The rock weighs three lbs.; and it is in virtue of this fact about the rock that certain events *involving* the rock have certain effects. But the rock *per se* is not a cause.

Nor is the rock an effect. Obviously the rock's existence is the end-result of a series of events. But it is these events, not the rock *per se*, that are appropriately described as 'effects.'

Notice that, while we would hesitate to describe the rock as an *effect* of the various events that led to its existence, we would not hesitate to describe the rock as the *product* of these events.

Conditions identical with event-patterns

Conditions, e.g. Smith's having high blood pressure, can be causes and effects. This is because conditions are event-pluralities and therefore events in their own right.

But not all event-pluralities are conditions: a condition is a stable pattern of events.

For your metabolic condition to remain unchanged is for the various events constitutive of it to continue to occur within certain limits; it is for your heart to continue to beat at a certain rate; it is for the cells composing your body to be duly oxygenated at regular intervals; etc. Unless you are cryogenically frozen, for your metabolic condition to remain unchanged is *not* for you to undergo *no* metabolic changes.

For the room's temperature to remain constant is *not* for the various particles involved to undergo *no* change. It is for the changes that they are constantly undergoing to remain within certain limits; it is for those changes to remain such that their mean kinetic energy doesn't change. Unless their mean kinetic energy is zero---and even then---this means that those air molecules are constantly whizzing about.

For there to be a change in the room's temperature is for there to be a change in the pattern of change constitutive of its having a given temperature. Supposing that the room's temperature at a given time is 78^{0}, for the room's temperature to increase is for there to an increase in the frequency of the various particle-collisions constitutive of the room's being 78^{0}; and for the room's temperature to drop is for there to be a decrease in the frequency of such collisions. The same thing *mutatis mutandis* holds with respect to changes in your physiological condition.

So-called changes typically identical with meta-changes

An event is a change. Therefore, all causes and all effects are changes. But we take for granted, and thus cease to be aware of, the occurrence of the changes constitutive of stable conditions; and we thus tend to reserve the use of the word ‘change’ for changes in patterns of change.

If the patient’s heart-rate, blood-oxygen level, etc. are stable, the attending physician will tell his supervisor that ‘there are no changes to report', his meaning that, whatever the patient’s condition was, the changes constitutive of his being in that condition have remained within certain limits.

An economist says that ‘there has been no change’ in the economy if the patterns of buying and selling constitutive of a nation’s economy haven't changed.

An accurate thermometer registers no change in the room’s temperature if there has been no temperature-relevant change in the pattern of particle collisions involved in the room’s having that temperature.

Explanation, prediction, and control

Scientific knowledge enables us to explain past events and to predict future ones, and it thereby enables us, within limits, to control the course of events. Supposing that X is some miracle-cure for some ailment that you have, you know how to improve your health-condition (*you should take X*); you can explain why your condition improved (*you took X*); and you can explain why your condition deteriorated (*you stopped taking X*).

Stable-event patterns a prerequisite for explanation-identification

In contexts in which there are stable patterns of events, it tends to be possible to explain, predict, and control what happens; and it tends to be *im*possible to do so in contexts in which there are no such patterns and which therefore consist of rogue events. It is impossible to predict what will happen to a 10 MG tablet of X that is being blown about by heavy winds; it is relatively easy to predict what will happen such a tablet is consumed by someone with the aforementioned ailment.

It is for this reason that explanation and prediction almost always consist in the identification of *program*-causes. Given some event E_n, E_n's *program*-cause is a structure or stable condition that, although not *by itself* determinative of E_n, predetermines the occurrence of some structure-internal specific event E_{n-1} that *is* thus determinative.

The *immediate* cause of Smith's getting mugged is Brown's desire for money. The program-cause is the disastrous financial condition of Smith's country: since that country is poor, it is inevitable that *somebody* will violate *somebody's* property-rights in *some way or other*. The economist is right to blame the mugging on the evils of capitalism. The psychologist is right to blame the mugging on Brown's lack of scruples. The economist has correctly identified the *program*-cause of Smith's mugging; the psychologist has correctly identified the *immediate* cause.

Freud on causation in relation to the development of mental illness

Freud correctly says that, if a given person X succumbs to a given psychology malady M, four different kinds of causes are likely to be operative:

(1) X's precondition (how X was before he was afflicted with M);
(2) M's specific cause (the trauma that, though it did not itself suffice for X's succumbing to M, made it inevitable that the slightest disturbance in X's condition would thus suffice);
(3) M's concurrent cause (some trauma experienced by X, e.g. a recent break-up, that is operative independently of (2) and isn't constitutive of X's precondition);
(4) The immediate or precipitating cause of M (some specific event, e.g. the sound of a car backfiring, that pushes X, who has thus far been on the threshold that separates those who suffer from X from those who don't, just past that threshold).

Psychological overdetermination

Freud says that mental events are often ‘overdetermined.’ What does Freud mean?

The obvious answer is: A given event, e.g. x’s falling in love with y, is overdetermined iff that event has multiple *independently* operative sufficient causes.

An event’s immediate cause is not independent of its program-cause; nor is it independent of any one of its non-immediate causes.

An example of this form of overdetermination: x falls in love with y because of y’s intelligence as well as y’s looks; and each of these factors was so powerful that, had the other one been altogether absent, it would by itself have sufficed for x’s falling in love with y.

But even though this is *a* form of causal overdetermination, it isn't the kind Freud has in mind. Suppose, as before, that x falls in love with y, both because of y’s looks and because of y’s intelligence, *but* that, were x intelligent but not good-looking or *vice versa*, x would not have fallen in love with y. Under those circumstances, Freud would *still* say that x’s falling in love with y was overdetermined.

When Freud says that a given psychological event e is ‘overdetermined,’ he means that it was brought about by the confluence of two or more psychological events each of which *justified* e’s occurrence. Relative to any one such event, e's occurrence made sense, even though not a single one of those events was necessarily capable on its own of bringing about e.

Thus, if Freud was right, as he surely was, causation in the psychological sphere often involves the exceeding of *thresholds*: it isn’t enough that x have just one reason for falling in love with y; unless he has several reasons, nothing happens; but once enough such reasons accumulate, the damn bursts and x falls in love with y.

Threshold effects and the analog-digital distinction

Discontinuous changes tend to involve threshold effects. A circuit breaker doesn't do anything until there is electricity-overload, at which point it shuts off the electrical pathway in question. A firm doesn't declare bankruptcy until it's in the red by a certain amount, at which point it files the paper work. Continuous changes don't involve threshold effects.

To the extent that the world changes discontinuously, it has a *digital* structure, and to the extent that it changes continuously, it has an *analogue* structure. An analogue clock is one that changes continuously with the passage of time; a digital clock is one that abruptly changes states at predetermined intervals.

Causes as INUS-conditions **("INUS"=insufficient, non-redundant, unnecessary, sufficient)**

Suppose that

e (Smith's flipping the switch)

is the cause of

e (the light's turning on),*

but that e* was not overdetermined e*, i.e. that e* wouldn't have happened unless e had happened.

By itself, e obviously isn't *sufficient* for e*. If the bulbs are broken or the circuitry is damaged, e will not bring about e*.

But given that e* did happen, it follows that there was *some* confluence of events/conditions that sufficed for e*; and given that under the operative circumstances e* would not have happened unless e had happened, it follows that e is a *non-redundant* part of that confluence of events.

That confluence of events, though sufficient for e*, wasn't *necessary* for e*: after all, an entirely different confluence of events---one not involving my flipping the switch or any of the events intervening between my doing so and the light's turning on --- could have caused the light to turn on.

It follows that e was an *insufficient but non-redundant part of a condition that unnecessary but sufficient* for e*'s occurrence --that, in other words, that e was an *INUS condition* for e*. In most context[3],

(i) Such and such was the cause of thus and such

means

(ii) Such and such was an INUS condition for thus and such.

Counterfactual conditionals

Statements of the form if P had been the case, then Q would have been the case are *counterfactual conditionals* (or *counterfactuals*). Thus, the meaning of

(1) If Kennedy hadn't been assassinated, the US would have pulled out of Viet Nam before 1965

is

(1) Kennedy's assassination prevented a pre-1965 troop-withdrawal.*

Similarly, the meaning of

(2) If Smith hadn't fall out of the tree, he wouldn't have broken his leg

is

(2) Smith's fall caused him to break his leg.*

Thus, many counterfactuals are to the effect that one thing caused another to happen and many counterfactuals therefore hold in virtue of dependence-relations of the causal variety.

But some counterfactuals hold in virtue of dependence-relations, not of the causal, but of the logical variety. For example, the meaning of

(3) If, at time t, Smith had had over $10, then, at time t, he would have had over $5,

is

(3) Smith had over $10 at t' entails 'Smith had over $5 at t.*

Thus, all counterfactuals affirm the existence of instances of necessitation of some kind or other, and all therefore counterfactuals therefore hold in virtue of dependence-relations of one kind or other.

In this section, we will focus exclusively on counterfactuals of the causal variety.

David Lewis's analysis of counterfactuals

According to David Lewis (1972, 1984), the meaning of (1) is

(1#) If w is a world in which Kennedy was not assassinated but is otherwise as much like our world as that fact allows it to be, then in w the US pulled out of Viet Nam before 1965.

Lewis's analysis is unacceptable. Any given event has deep causal roots. Kennedy's assassination was the consequence of a confluence of political, psychological, and physical events, and each of those events is in its turn a consequence of some other confluence of events, and so on. This means that if Kennedy's assassination doesn't take place in a given world, then either

(a) That world differs so much from our world that it is no guarantee that Kennedy even exists in it, let alone effectuates an early troop-withdrawal

or

(b) That world differs from ours in respect of the deterministic mechanisms involved in it,

Given either (a) or (b), there is little or no chance that, in the world in question, Kennedy effectuates an early troop-withdrawal. So Lewis' analysis is false.

Also, it is by examining the causal and logical structure of our world, not other worlds, that we know whether or not a given counterfactual is true. This means that counterfactuals hold in virtue of the logical and causal structure of our world, as opposed to other worlds.

Since there is no reason, apart from Lewis's analysis, to believe that there are other worlds, it follows that there is no reason at all to posit their existence or therefore to accept Lewis's analysis.

Program-causality

Because the water in the vessel is boiling, water-molecules strike the interior walls of the vessel with so much force that the vessel cracks. Let $c_1 \ldots c_n$ be those high-speed collisions. The water's boiling did nothing in the way of cracking the vessel except in so far as it was responsible for $c_1 \ldots c_n$.

Nonetheless, for reasons having nothing to do with our not knowing about those specific particle-collisions, the appropriate answer in some contexts to the question "why did the walls of the vessel crack?" is "because the water inside the vessel was boiling," and *not* "because $c_1 \ldots c_n$ occurred." *Ceteris paribus* the vessel would have cracked even if $c_1 \ldots c_n$ hadn't occurred, since the conditions in question guaranteed the occurrence of *some* such particle-collisions.

Given that the cause of an event is what guarantees its occurrence, the boiling of the water is, insofar as it guaranteed the occurrence of the likes of $c_1 \ldots c_n$, appropriately described as "the cause" of the cracking of the vessel.

Causal redundancy in relation to counterfactuals

(i) $c_1 \ldots c_n$ caused the vessel to crack

is true. But

(i) if $c_1 \ldots c_n$ had not occurred, then the vessel would not have cracked*

is false. If $c_1 \ldots c_n$ had not occurred, other particle-collisions *would* have occurred that would have cracked the vessel.

In general, where there is causal redundancy, causal truths cannot be converted into counterfactual truths. Suppose it a truth that

(A) The bird died because Smith shot it.

If the bird was otherwise unthreatened, then a consequence of (A) is

(A) If Smith hadn't shot it, the bird would not have died.*

But if Smith was part of a pack of hunters, each of whom would have shot the bird had Smith not done so, then (A*) is not a consequence of (A).

The Counterfactual Analysis of Causality (CAC)

According to CAC, the *meaning* of

*(#) e is the cause of e**

is

(#) If e hadn't happened, then e* wouldn't have happened.*

The motivation for CAC

Ordinarily the meaning of (#) is taken to be

(#^) e compelled e to occur, i.e. e's occurrence necessitated e*'s occurrence.*

Some philosophers and some physicists have argued that there is no empirical (observational) evidence of the existence of forces. This being their position, they have two options.

Option #1: Take the deeply revisionist position that there is no observational support for any causal claim.

Option #2: Hold onto the presumption that causal claims are often observationally well-supported but reinterpret those claims in such a way that their veridicality isn't predicated on the existence of (supposedly) undiscovered forces.

CAC is meant to provide one such way of reinterpreting causal claims.

CAC evaluated

Two questions must be asked in connection with CAC? First, is it *extensionally* correct? Second, is it *intensionally* correct?

A theory is *extensionally* correct if there are no counterexamples to it.

A theory is *intensionally* correct if, first, there are no counterexamples to it and, second, the reason there are no counterexamples to it is that it is a correct description of what it's supposed to describe.

Intensional correctness entails extensional correctness, but not *vice versa.*

Suppose I put forth the following theory: *There will be at least 10 articles of furniture in Olin 310 on Thursday, and the reason for this is that the invisible and otherwise undetectable Martian who lives in my closet will telekinetically transport at least 10 desks into Olin 310*.

Given that, as it happens, there *did* turn out to be more than 10 articles of furniture in that place at that time, this theory is *extensionally* correct, meaning that there are no counterexamples to it; but it is *intensionally* incorrect, meaning that, even though the facts are consistent with it, they are consistent with it for the wrong reason.

CAC is extensionally false. For

(#) If e hadn't happened, then e* wouldn't have happened*

is not always consequence of

(#) e is the cause of e.*

(#*) fails to be such a consequence whenever e is but one of many would-be causes of e*, and it therefore fails whenever circumstances *program for* e*'s occurrence.

CAC is also *intensionally* false. Suppose it true both that

(A) The otherwise unthreatened bird died because Smith shot it

and also consequently that

(A) if Smith hadn't shot it, the bird would not have died.*

If it weren't for (A) or for some other causal claim, there would be no justification for (A*).

Another example:

(I) Because the pancake batter was poured onto the hot griddle, the batter turned into a pancake.

Supposing that no other person or force would have poured the batter onto the griddle, a consequence of (I) is

(I) If the pancake batter hadn't been poured onto the griddle, it wouldn't have turned into a pancake.*

But whereas one can reasonably accept (I*) on the basis of (I), one cannot possibly accept (I) on the basis of (I*): knowledge of (I*) is necessarily derivative of knowledge of (I), suggesting that the truth affirmed by (I) is more fundamental than the truth affirmed by (I*).

This suggestion is consistent with the fact that, whereas (I*) is a consequence of (I), (I) is not a consequence of (I*). *Explanation*: Given only that the batter would not have turned into a pancake had it not been poured onto the griddle, it follows that its being poured on the griddle was in some way or other causative of its turning into a pancake. But given only that the batter turned into a pancake because it was poured on the griddle, it does *not* follow that the batter wouldn't have turned into a pancake had it not been poured on the griddle. If the batter was in a pancake shaped vessel which suddenly heated up, then it would have turned into a pancake without being poured onto a griddle.

CAC Defended (CACD)

Here is the one and only way in which advocates of CAC can defend their position against the just-made points: *Once again suppose that*

(A) the bird died because Smith shot it.

If Smith missed and Smith's hunting partner Brown took the opportunity to shoot the bird dead, the effect of Brown's shooting the bird would have been different from the effect of Smith's doing so. Depending on whether it is Smith or Brown who shoots the bird, the bird dies at this as opposed to that time; the bullet strikes the bird at this as opposed to that angle; the slain bird describes this as opposed to that trajectory as it falls to the earth, where it lands in this as opposed to that location, at this opposed to that exact moment in time.
In general, supposing that e is the cause of e, it is only if e is insufficiently described that*

(#) if e hadn't happened, then e* wouldn't have happened*

fails to be a consequence of

(#) e is the cause of e.*

Problems with CACD

There are three problems with CACD. These have to do with

(i) Events whose causes involve the exceeding of thresholds,
(ii) Events whose causes are mechanisms that can be implemented in different ways, and
(iii) Events whose causes program for their occurrence.

(i) Explained

At time t, I push the button and the elevator comes as a result. Had I not done so, someone else might have pushed the button at t and the elevator would still have come, and the result would have been the same: the elevator would have come at the same time, in the same way.

Of course, there is little chance that this other person would have pushed the button in the exact way that I did: he would have used either more or less force than I did, and his finger wouldn't have come into contact with the exact same parts of the button as my finger.

But this is irrelevant. The button is a binary-mechanism. When the elevator comes isn't a function of the way in which the button is pushed, and neither is the manner in which the elevator comes. Thus, had some other person pushed the button, the result would have been the *very same event* that resulted from my pushing the button.

Therefore CACD is false, given that, according to it, the (hypothetical) event resulting from the other person's pushing the button would *not* have been numerically identical with the (actual) event resulting from my pushing it. In general, the identity of a given threshold effect is not contingent on the specific composition of its antecedents. Thus, where threshold effects are concerned, there *is* genuine causal redundancy, and CACD is therefore false.

(ii) explained

I type "1+2=" into a calculator; then I hit ENTER. My hitting ENTER caused a "3" to appear on the LCD. But if I hadn't hit ENTER, the "3" still would have appeared, since my friend Gus would have hit ENTER had I not done so.

Had Gus hit ENTER, he would almost certainly not have hit it with the same degree of force or in the same manner, a likely result being that the electrical activity precipitated by his hitting ENTER would have differed in some respects from the electrical activity precipitated by my hitting ENTER. But so long as Gus didn't strike the ENTER key with excessive force or otherwise cause the calculator to malfunction, those differences wouldn't have caused a different operation to be carried out (they wouldn't have caused 50 to be divided by 5); nor would they have caused the same operation to be carried out in a different way (they wouldn't have caused the computer to add them in base 3, as opposed to base 10 or base 2).

Thus, the differences between the Gus-caused electrical activity and its JM-caused counterpart led to no differences between the Gus-caused software activity and its JM-counterpart. Since the Gus-caused software activity did not otherwise differ from its JM-caused counterpart, the former is numerically identical with the latter. Since the Gus-caused electrical activity *did* differ from its JM-caused counterpart. And since the JM-caused software activity was identical with its Gus-caused counterpart, that software-activity was in each case distinct from the electrical activity that mediated it.

In general, software-level activity---meaning, ultimately, any causal-series that can be implemented in different ways---is to be distinguished from the event-series by which it is implemented, notwithstanding that the two series are spatiotemporally coincident. In general, the identity an event that can be implemented in different ways isn't contingent on its exact composition or, therefore, on the exact composition of its cause. Thus, where such events are concerned, there *is* genuine causal redundancy, and CACD is therefore false.

(iii) Explained

A hot griddle converts liquid pancake-batter into a solid pancake. Let $m_1 \ldots m_n$ be the confluence of micro-events that is immediately responsible for this conversion process.

Given that the batter was on a hot griddle, it was inevitable that *some* such micro-events occur, even though it obviously wasn't inevitable that those specific micro-events occur. Therefore, if $m_1 \ldots m_n$ hadn't occurred, the conversion-process still would have occurred, thanks to some entirely distinct confluence of micro-events $m^*_1 \ldots m^*_m$.

The conversion process precipitated by $m_1 \ldots m_n$ has a different internal composition from the hypothetical conversion process precipitated by $m^*_1 \ldots m^*_m$.

But just as the identity of a military operation isn't contingent on the exact identities of the infantrymen who are carrying out the general's orders, so the identity of the conversion-process in question isn't contingent on the identities of the exact micro-particles that happened to be immediately responsible for it, a consequence being that the hypothetical conversion-process is numerically identical with the actual conversion-process.

Therefore CACD is false, given that, according to it, the (hypothetical) conversion-process would *not* have been numerically identical with the (actual) conversion-process. In general, the identity of an event whose cause programs for it is not contingent on its specific composition or on the precise identity of its immediate antecedents. Thus, where such events are concerned, there *is* genuine causal redundancy, and CACD is therefore false.

Summary of our discussion of The Counterfactual Analysis of Causality (CAC)

According to CAC, e is the cause of e* iff e* wouldn't have happened had e not happened.

Frankfurt's proves that CAC is false. His argument: *My pushing the button is the cause of the elevator's coming, but the elevator might have come even if I hadn't pushed the button—somebody else might have pushed it. Therefore, CAC is false.*

Attempts to falsify Frankfurt's critique are false. Therefore CAC is demonstrably false, and not merely *prima facie* so.

Chapter 4 The Regularity Analysis of Causation

The Regularity Analysis of Causation stated

According to Hume, e_1 is the direct cause of e_2 iff any event that is similar to e_1 spatially adjacent to an immediately posterior event that is similar to e_2; and e_1 is the indirect cause of e_n iff, for each i ($1 \leq i \leq n$), any event that is similar to e_i spatially adjacent to an immediately posterior event that is similar to e_{i+1}.

Hume holds that, if e is the direct cause of e*, then there is a law to the effect that any event similar to e is adjacent to an immediately posterior event that is similar to e*.

Hume therefore holds that laws are regularities and causal sequences are instances of such regularities.

Hume therefore rejects the conventional view, according to which laws are structures that are responsible for such regularities and are therefore distinct from them.

Hume therefore rejects the conventional view, according to which causes are responsible for their effects and are therefore not merely their antecedents.

Thus, in Hume's view, laws are regularities and causal-series are instances of regularities.

Hume's argument for this analysis (first version)

Says Hume (I am paraphrasing): *If there is any empirical (observational) basis for saying that e caused e*, it is*

(i) The fact that e was e's immediate predecessor and neighbor,*

coupled with

(ii) The fact that events similar to e are always the immediate predecessors and neighbors of events similar to e.*

Hume's argument for this view (second version)

Let e and e be two events such that:*

e is identical with A's colliding with hitherto stationary B,

and such that

e is identical with B's thereupon begin to roll.*

And let e-e be the sequence consisting of e and e*, in that order.*
e-e is a paradigm-case of a causal sequence.*
Let us suppose that each of e and e was observed by maximally good vision, whose observations were therefore maximally precise.*
It is presumed that e compelled e to occur and that its doing so involved more than its immediately preceding and being adjacent to e*.*
But no such compulsion was observed. This is not the relevant sense-perceptions were imprecise. It is because there are no forces to observe.

The fallacy embodied in Hume's analysis—Part 1

Hume's analysis of causality, and his defense of it, embody a failure to grasp some of the most basic truths about the causality. In this section (Part 1), we will state those truths. In the next (Part 2), we will explain why they warrant a rejection of Hume's analysis.

There are two kinds of causal processes:

(i) Those that involve interactions

and

(ii) Those that consist solely of instances of persistence.

An example of (i) would be a projectile (e.g. a tennis ball, a light beam) that bounces off of a wall.

An example of (ii) would be a projectile that doesn't collide with anything or otherwise undergo motion-altering interactions with other objects. Such a projectile's state of motion at time t is a causal consequence of its pre-t state of motion; the projectile's mass at t is a causal consequence of its pre-t mass; any given fact about the projectile's physical condition at t is the effect of the projectile's pre-t condition.

Following Bertrand Russell[4], we will refer to type-(i) causation as "transeunt" causation and to type-(ii) causation as "immanent" causation.

Anything that endures for any length of time is a series of instances of immanent causation and is thus a causal series: a series of states of affairs such that any given such state, apart from the last, is the cause of a structurally similar such state and any such state, apart from the first, is the effect of a similar such state. The occupants of space-time are causal series.

The difference between immanent and transeunt causation is that immanent causation is, whereas transeunt causation is not, structure-preserving. The stain-glass window in the church is a structure-preserving causal series; so is the vase that was just flung at it. When the vase collides with the window, these two causal series converge, and the result is a causal series that is *not* structure-preserving. The vase and the window shattered upon impact; so their respective structures were thereby destroyed.

Any given case of transeunt causation involves the cessation of a structure-preserving causal series. Ball A collides with hitherto stationary Ball B; A stops moving and B starts moving. Even though each ball has retained its structure, neither ball's *condition* has retained its structure. By virtue of immobilizing A and mobilizing B, the collision introduced a sudden, discontinuous change in the nature of causal-series constitutive of their respective existence.

The operative words are "sudden" and "discontinuous." Objects that are not acted upon sometimes disintegrate and therefore lose their structure. But such structural changes are continuous. A melting ice-sculpture loses its structure in a continuous manner, its loss of structure being a consequence of immanent causation. An ice-sculpture that is smashed by a vandal loses its structure in a discontinuous manner, its loss of structure being a consequence of transeunt causation.

There are obviously incalculably many instances of transeunt causation. But the *relative* number of such instances is zero. (A real number can be zero even if the corresponding cardinal is not zero.) This is because any instance of transeunt causation is a *singularity*---a destroyer of

the stable patterns of causal interdependencies that constitute the occupants of the space-time manifold. Any case of transeunt causality is a breakdown of causal integrity. Transeunt causation involves not only involves a weakening, but a severing, of causal ties. The essence of transeunt causation is the discontinuous alteration of object-constitutive causal-series.

The fallacy embodied in Hume's analysis—Part 2

The essence of transeunt causation is discontinuity. The essence of immanent causation is continuity. A consequence is that the world has a digital structure to the extent that it hosts instances of transeunt causation, and an analogue structure to the extent that it hosts instances of immanent causation. A chain-reaction, consists of discrete events. (The first domino falls; then the second one falls; and so on.) The movement of a projectile doesn't have such a structure; nor does the continued existence of a stationary object, such as a desk-weight. Neither such series consists of discrete, successive states of affairs. *A fortiori* neither series consists of entities whose causal interrelations to each other have to be explained. All instances of continuity are causal series, but no such instance consists of discrete states whose causal ties to their successors need to be identified.

Causal explanations are necessary only when they're possible, and they're possible only when continuities have been disrupted. Hume takes it for granted that all instances of causation are identical with such disruptions, and Hume therefore takes it for granted that those continuities *per se* are non-causal. When trying to account for an instance of transeunt causation, the correct approach is to try to find what has been preserved. (When two balls collide, kinetic energy is preserved.) The correct approach is to look for a continuity underneath the discontinuities. (It's precisely because A's kinetic energy continues to flow continuously that each ball's state of motion changes.)

Because Hume took *breakdowns* of causal cohesiveness as *instances* of such cohesiveness, he believed causation to be a relation between disjoint entities. This in turn led him to misanalyse the instances of causation on which he focused: it led him to focus on the ways in which the effect was disjoint from the cause and to overlook the ways in which the effect was continuous with the cause.

The fallacy embodied in Hume's analysis—Part 3: The Concept of a Hume-world and its uselessness in the way of supporting Hume's analysis

Hume claims that there is no empirical basis for preferring

(a) e made e happen*

to

(b) e-similar events always follow immediately antecedent neighboring e-similar events.*

Hume's argument for this claim

An event exactly like e could immediately precede and be spatially contiguous with an event that it did not cause but was otherwise exactly like e*.

For example: Magnets underneath the surface on which A is rolling cause it to stop at the exact instant in which it makes contact with B, whereupon those same magnets cause B to start rolling. Further, those magnets change the states of motion of A and B in exactly the way in which, in actuality, the collision changed their states of motion. There is no observable difference between this hypothetical situation, on the one hand, and the actual situation, on the other.

What is more, in the hypothetical situation, the collision was had nothing to do with the subsequent changes in the balls' respective states of motion, and there was therefore no causal nexus between the collision and the motion-changes to be observed.

Given that the actual and the hypothetical situations are observationally indistinguishable, it follows that no causal nexus linking cause and effect can be observed in the actual situation. Given any instance of causation, a similar argument shows that it is observationally indistinguishable from a non-causal event-sequence. Thus, even though the putative causes of events can obviously be observed, their *acting* as causes cannot be observed.

Hume worlds observationally different from their non-Humean counterparts

But the hypothetical scenario in question *is* observably different, in many respects, from the actual one. If magnets are responsible for relevant kinetic changes, there will be observable evidence of that fact. That evidence may not be available on the occasion on which you witness the balls collide, but that is irrelevant. If A stopped for no reason at all---if it was sheer coincidence that its coming into contact with B happened at the same time as its ceasing to roll---the situation will be observably very different from one in which its ceasing to roll was a consequence of its coming into contact with B. In the one case but not the other, A and B will be deformed in certain ways, they'll experience changes in temperature, their centers of gravity will change, and so on. These differences won't be observable to the naked eye, but that means only that Hume's analysis of the situation holds, if at all, only with respect to an incomplete description of it.

In general, genuine causal processes *are* observationally distinguishable from pseudo-causal processes. And what observation reveals is that the events composing a genuine causal sequence are continuous with each other, whereas those events composing a pseudo-causal sequence are merely contiguous with each other, effects are *continuous* with their causes, and not merely contiguous with them.

In the scenario in which the post-collision kinetic changes undergone by A and B are the result of magnet-involving trickery, there is a continuous process of which B's moving is a segment and there is also a continuous process of which A's ceasing to move is a segment, but there is no continuous process of which both are segments. But there *is* such a process in the case in which the collision is in fact responsible for the changes undergone by A and B: those changes mediate a continuous transfer of momentum (mass x acceleration), much as changes in individual water-molecules mediate the rippling of water.

These points suggest that the essence of causation is continuity: a continuous event-series is a causal process; a discontinuous event-series is a non-causal process, even if every non-terminal member of that sequence is contiguous with its success and every non-initial member is contiguous with its predecessor. If this is correct, then whenever we observe a process to be continuous, we are *ipso facto* observing states of affairs to be causally responsible for other states of affairs.

Three additional objections to the Regularity Analysis of Causation

Problem #1: The temporal order is to be understood in terms of the causal order, not *vice versa*. Hume's analysis assumes otherwise and is therefore viciously circular and therefore false.

Problem #2: Correlation is not causation. A consequence of Hume's analysis---indeed, the very essence of it---is that correlation *is* causation. Therefore, Hume's analysis is false.

Problem #3: Given any spatiotemporally adjacent events, e and e*, there is some regularity of which that event-sequence is an instance. Hume's analysis therefore entails that any two spatiotemporally adjacent events are related as cause and effect. Since this isn't the case, Hume's analysis is incorrect.

Explanation of Problem #1

There is no way to assign relative dates or positions to events except on the basis of their causal interrelations. e is simultaneous with e* iff there is no causal process that begins with the one and ends with the other; and precedes e* iff either

(i) e affects e (e.g. there is a causal process, such as a light ray, that begins with e and ends with e*)*

or

(ii) There is some event that is simultaneous with e that affects e.*

e is spatially adjacent with e* iff either

(a) They are non-simultaneous and the one can instantaneously affect the other

or

(b) They are simultaneous, and therefore cannot affect each other, but there is some event e# such that each cannot instantaneously affect e#.

Newton said that the concept of instantaneous "action at a distance" is incoherent. He was right. If there is no time-delay between cause and effect, there is no distance between cause and effect, the reason being that for there to be either a spatial or temporal gap between cause and effect *is* for that instance of causation not to be instantaneous.

Statements about spatiotemporal interrelations have neither operational nor philosophical significance unless considered as statements about causal interrelations:

(1) e precedes e, but e doesn't affect e* or anything simultaneous therewith,*
which means

(2) e is no part of the history of the time-period occupied by e,*

which means

(3) e does not precede e,*

which means

(4) e and e are simultaneous occupants of some one spatial manifold, but neither can affect anything affected by the other,*

which means

(5) At no point in time is there any region of space whose history or whose present includes both e and e,*

which means

(6) e and e do not simultaneously occupy some one spatial manifold.*

Truths about spatiotemporal interrelations are therefore truths about causal interrelations, and the spatiotemporal order is to be understood in terms of the causal order.

Hume assumes otherwise, and his analysis is therefore incorrect.

Explanation of Problem #2

Ball A collides with Ball B and B moves as a result. This is caught on film; a movie is made out of this footage; the movie consists of endless repetitions of the same footage---of A striking and thereby displacing B. In this scenario, whenever there is an *image* of A striking B, there is a subsequent, spatiotemporally adjacent *image* of B moving. But the A-*image* isn't for the subsequent B-image: there is a projector that causes there to be an A-image and also causes there to be an immediately subsequent, adjacent B-image. This projector is responsible for each image. The A-image *per se* is no more capable of producing a B-image than a shadow of one man stabbing another can be responsible for the second man's death. But a consequence of RA is that image-A *is* responsible for image-B. Therefore, RA is incorrect.

The sequence consisting of those two images is an instance of *epiphenomenal causation*. Epiphenomenal causation is not actual causation; it is pseudo-causation. e is the epiphenomenal cause of e* iff

(i) Any e-similar event is immediately succeeded by an adjacent e-similar event,*
(ii) e is not the cause of e, and*
(iii) The cause of each of e and e is some underlying causal mechanism e#.*

Reality abounds with instances of epiphenomenal causation. If a thermometer is immersed in hot water, the upward-movement of the column of mercury immediately precedes an adjacent acceleration of the particles composing the walls of the thermometer. But the former is not the cause of the latter; the cause of each is the fact that the thermometer in hot water. The fact that the two events are spatiotemporally adjacent has to do with the structure of the thermometer, coupled with the fact that different substances have different coefficients of expansion.

When my computer is powering on, a moving dot-circle appears on the monitor: each dot appears immediately before an adjacent dot and the series of dot-occurrences forms a circle. No two of those occurrences are related as cause and effect; each dot is created by invisible software activity that also creates a neighboring dot. According to RA, each dot *is* the cause of its immediate successor. A dot-occurrence in a given place is always immediately succeeded by an adjacent dot-occurrence. Therefore, RA entails that neighboring dots are related as cause and effect. Therefore, RA is incorrect.

Computers are structured in such a way that *bona fide* causal processes are mirrored by pseudo-causal causal processes. Visible desktops operations are manifestations of invisible software activity. Software events that are related as cause and effect create desktop-events that appear to be related as cause and effect but are not so related. These desktop-occurrences are visible representations of otherwise invisible software activity, and because we can be control these desktop-occurrences, by controlling the icons involved therein, we are able to control the aforementioned invisible, and otherwise uncontrollable, software activity.

In some cases, a given desktop-occurrence is the *indirect* cause of another such occurrence. Moving the cursor from one place to another causes it to change shape: the location-change precipitates electrical activity that brings about a shape-change. But even this sort of event-sequence is inconsistent with RA. Supposing that a given change in the cursor's location is *immediately* followed by a change in its shape, the location-change cannot possibly be the *indirect* cause of the shape-change: nothing can possibly be the *indirect* cause of anything that *immediately* follows it. And since the location-change obviously isn't *directly* causally

responsible for the shape-change, there exists *no* causal sequence of which both events are members. A *prior* location-change on the part of the cursor precipitated the electrical activity that brought about the just-mentioned location-change *as well as* the subsequent shape-change.

Explanation of Problem #3

Consider the previously mentioned dot-circle. Let e be one of those dots; let e* be the dot that immediately follows e; and let r and r* be the respective spatiotemporal locations of those occurrences.

If k is the class of all dots occurring in r, and k* is the class of all dots occurring in r*, then any given k-occurrence is immediately followed by a neighboring k*-occurrence.

Given any event E, there is some category K of which E is the sole instance. It follows that, given any two neighboring events, there is some regularity of which that event-sequence is an instance. According to RA, any such sequence is *ipso facto* causal. Since this isn't the case, RA is incorrect.

Chapter 5 Causation to be understood in terms of the shortcomings of the deductive-nomological model of scientific explanation

The deductive-nomological model of explanation (DN)

According to DN, E is explained iff, supposing that C is a description of the circumstances in which E occurred and L is a description of the laws operative in C, then the proposition that E occurred is a logical consequence of the conjunction of C and L.

Here are four examples of DN-compliant explanations:

Example #1:
E: the metal bar expands at t.*
C: the metal bar is heated at previous time t.
L: metal expands when heated.

C+L (C conjoined with L) logically entails E. Therefore, according to DN, the explanation of E is C+L.

Note: This argument presupposes that the time-lag between t and t* doesn't scuttle the entailment; it presupposes that nothing happens during that interval that doesn't stymie the metal bar's expansion. Given that there, for any distinct time-intervals, there *is* a time-lag between them, this requirement isn't always met. So, technically, C+L does *not* entail E. Rather, C+L entails E *on the condition* that nothing has changed during the time that has elapsed between t and t*. This is a problem for DN and for RA. But, in the context of this handout, we will assume that the time-lag between t and t* doesn't make a difference and that the entailment goes through. This fact about Example #1 is true *mutatis mutandis* of each of the following illustrations of DN.

Example #2:
E: x is sinking at time t.*
C_1: x was placed in body of water y at t.
C_2: x's density exceeds y's.
L_1: a solid body floating in water (or liquid of any kind) displaces a volume of water whose weight equals that of the solid body.
L_2 (a corollary of L_1): A solid object floats/sinks in liquid b exactly if a's density is less than/greater than b's.

E is a logical consequence of—and, according to DN, is therefore explained by---this statement-set.

Example #3
E: The rate at which x is approaching y is increasing exponentially.
C_1: x's mass is m and y's mass is m.*
C_2: x and y jointly constitute an isolated system.
L_1: the gravitational attraction between two bodies is directly proportional to their masses and inversely proportional to the square of the distance between them.
L_2 (a corollary of L_1): ceteris paribus, the gravitational attraction between two bodies not only increases, but does so exponentially, as they approach each other.
L_3: (a corollary of L_2): ceteris paribus, the rate at which two bodies approach each other not only increases, but does so exponentially, as they approach each other.

E is a logical consequence of---and, according to DN, is therefore explained by—this

statement-set.

Example #4:

E: the rate at which x is moving around is greater at t than it was at previous time t. (x is a piece of pollen on the surface of a body of water w.)*

C_1: w's temperature at t is greater than it was at t.*

C_2: all other factors have remained constant during that interval.

L_1: water consists of molecules.

L_2: the hotter a body of water, the more quickly the molecules composing it move.

L_3: the faster an object moves, the greater the force it exerts an object with which it collides, a consequence being that ceteris paribus the greater is the extent by which that collision displaces the impacted object.

E is a logical consequence of---and, according to DN, is therefore explained by—this statement-set.

Problems with DN

Laws are given by propositions that support counterfactuals. *Increasing gas-temperature increases gas-volume* is a law because, other things being equal, the greater the gas's temperature, the greater the volume it occupies. If the gas is confined to an unbreakable container, and when for that reason other things are *not* equal, then the gas's overall volume does not increase with temperature. What increases is the pressure of the gas on the container-walls. Hence the law: *pressure is directly proportional to temperature and inversely proportional to container-size.*

But this law is obviously a corollary of the first law that more temperature leads to more volume. For when the volume-increase is thwarted, it is replaced by a pressure on the gas's part to let its volume be increased.

Laws, so says the DN-advocate, are given by statements of the form *for any x, if x has phi, then x has psi*. In other words, laws (according to DN) are given by *universal, conditional generalizations.*

Nomic vs. accidental generalizations

But, as the DN-advocate recognizes, not all such universal generalizations expression laws. Cf. every coin in my pocket is a quarter (*for any x, if x is a coin in JM's pocket, then x is a quarter*). Not a law.

Why not? The obvious (and obviously correct) answer: a thing's being a coin in my pocket doesn't cause it to be a quarter or---to put it more accurately and also more generally---a thing's being a quarter doesn't *depend* on its being a quarter.

But the DN-advocate cannot accept this response. Suppose that the reason *metal expands when heated* is a law, whereas *all the coins in JM's pocket are all quarters*, is that heating metal causes it to expand, whereas putting coins in my pocket doesn't cause them become quarters. In that case, the fact that *metal expands when heated* is a law is contingent on the fact the fact that metal is already known to expand when heated. If this fact about metal is not already known, then there is no more reason to regard *metal expands when heated* as a nomic generalization and, therefore, as different from an obviously accidental generalization, such as *all the coins in JM's pocket are quarters.*

Why this is a problem for DN: Suppose it known, *before* it's known that *metal expands when heated* is a law, that the reason metal object A is expanding is that it's being heated. Suppose that, before this law is known, it's known for any a metal object M, it is known that it's because M is being heated that M expanding. In that case, the cause of M's expanding is known and so therefore is the explanation for the fact that it's expanding. (The cause: M's being heated. The explanation: M's being heated.) To know the explanation is to know the cause.

So in at least some cases, metal object x's expanding is not to be explained in terms of some covering law; it is *not* to be explained in terms of the law that metal expands when heated. *On the contrary, that law is to be explained in terms of the fact that x expand when heated.*

In general, if the difference between nomic and accidental generalizations is that the former have causal content whereas the latter do not, then DN is eviscerated; for the essence of DN is that events cannot possibly be explained, even approximately, except in terms of laws. And this *is* indeed the difference between nomic and accidental generalizations, and, consequently, DN *is* indeed eviscerated.

How DN-advocates deal with the distinction between nomic and accidental generalizations

For this reason, advocates DN do *not* hold the difference between *metal expands when heated* and *all the coins in JM's pocket are quarters* is that the former is a causal claim, whereas the latter is not. They therefore have to find some other relevant different difference between the two statements.

Hempel's proposal (widely accepted among DN-advocates, since they have no other choice) is this: whereas *metal expands when heated* supports counterfactuals, *all the coins in JM's pocket are quarters* does not. It is not the case that *if* a given non-quarter-coin were put in my pocket, it turn into a quarter. By contrast, it *is* the case that, if pinewood object x were placed in body of water y, it follows that x *would* float in y if x *were placed* in y. Thus, a law is a generalization that supports counterfactuals.

In this way, DN has seemingly insulated itself against the charge that nomic statements express causal relations and therefore against the consequent charge that DN is wrong to hold that explanations are categorically law-based and, in addition, against the charge that DN is fundamentally wrong to deny that nomic truths (truths that express laws) are discovered by consolidating various instances of causal knowledge.

DN viciously circular if it requires laws to support counterfactuals

Counterfactuals are cryptocausal statements; to know causes is to know explanations; to explain an event is to identify its cause. Therefore, DN is guilty of vicious circular in so far as its legitimacy is predicated on laws' being counterfactual supporting.

Explanation: Let L be the law the pinewood floats.[5] Why does L support counterfactuals? Because, if x is an arbitrary piece of pinewood, putting it in water *causes* it to float.

Another example: Metal expands when heated (*for any metal object x, if x is heated, then x expands*.) This law[6] supports counterfactuals. Why? Because, given some arbitrary metal object x, x's being heated *causes* x to expand.

Universal generalizations are statements to the effect that each member of a class of statements is true, each such member being a conditional proposition (one of the form *if P, then Q* or therefore, potentially, of the form *if x has phi, then x has psi*). If a given universal generalization is to support counterfactuals, and is thus, by Hempel's lights, to qualify as a law, then each of the corresponding conditionals must support counterfactuals.

But those counterfactuals support conditionals *only* if the following condition is met: Supposing that *if A has phi, then A has psi* is an arbitrary one of those conditionals, A's having phi is *causative* of A's having phi; i.e. A's having phi must *cause* A to have psi. If that condition isn't met, then if A *had phi, then A would have psi* isn't true, in which case *for any x, if x had phi, then has psi* is an accidental, not a nomic, generalization.

Further, if for all x, if x has phi, then x has psi is to *known* to be a law, then we must *know* that a given thing's having phi causes it to have psi. So, if we are to know that *for all x, if x has phi, then x has psi* is a law, then we must know that, given some arbitrary instance A of phi, A's being a phi is what is responsible for A's being a psi. But if we know *that*, then we *already* have knowledge of at least explanation (namely, that A's being a phi caused it to become a psi). And while it may be that our knowledge of *that* particular instance of causation is based on some other law L* (*if x is phi*, then x turns into a psi**), everything just said about the first law (*if x is phi, then x turns into a psi*) is equally true of L*---and it is therefore true of any law L** on which L* is based, and so on.

Since we must therefore eventually get to some law L# our knowledge of which is not derivative of our knowledge of some other law; and since L#, being counterfactual-supporting, is to the effect that, for arbitrary X, X's having PHI is causative of X's having PSI (for some properties PHI and PSI); and since, by virtue of knowing of L# and *ipso facto* knowing that X's having PHI is causative of X's having PSI, we *already* know how to explain at least one thing (namely, X's having PSI); it follows that explanations are *not*, at least not in general, law-based and, more precisely, that they do *not*, at least not in general, have the form *to explain E is to identify known laws* $L_1 \ldots L_n$ *such that, if* $C_1 \ldots C_m$ *is a description of the antecedents of E, then E is a logical consequence of* $L_1 \ldots L_n + C_1 \ldots C_m$.

Additional DN-internal vicious circularity

According to Hempel, if we want to know that the rope snapped, we need to know

(i) What the conditions were that obtained prior to the rope's snapping,

and

(ii) What the natural laws were in virtue of which those conditions led to the rope's snapping.

Some problems with this analysis

Let E be the rope-snapping in question. Supposing that E happened at time t, the term "the conditions that obtained prior to E" obviously doesn't refer to the state of the whole pre-t universe. Less obviously, it doesn't even refer to some small region the universe. Let R be a space-time region hosting the forces of which E was a consequence. Not everything in R was relevant to E. The optical conditions in R are irrelevant; the color of the object at the end of R is irrelevant; and so on.

So when Hempel tells us that E is to be explained in terms of "the conditions" preceding it, what Hempel means, assuming it true, is that the rope-snapping is a consequence of the *relevant* preconditions (e.g. the integrity of the rope, the weight of the suspended object, the distance of that object from the body that is dragging it downward, the density of that body).

If C is one of E's preconditions, C is a *relevant* precondition only if C *caused* E. [7] We can't provide a deductive-nomological explanation of E unless we know, for each E's infinitely many preconditions, whether or not that precondition is relevant to E; and if we know a given precondition C is relevant to E, then we *ipso facto* know that C caused E, and, since to explain an event is to identify its cause, we *ipso facto* have an explanation of E. Thus, in order to provide a deductive-nomological explanation of E, we must already have an explanation of E.

An objection: "Suppose you know that C caused E, but you don't know *how* C did so---you don't know what the operative laws were. In that case, your knowledge of what it is that caused E is incomplete."

Response: For me to have a *high-resolution* understanding of why E happened, I do indeed know the *how*, and not just the *what.* For me to know why Smith is sick it is not enough for me to know what made him sick (he ate rotten fish); I also need to know *how* his eating rotten fish led to his being ill. But if I know, as I surely can, that Smith is sick because he rotten fish, I surely do have knowledge of *an* explanation of his condition, even though that explanation is low-resolution.

Similarly, if I know *what* led to the rope's snapping (it was made of bubble gum and there was a heavy ball attached to it), I am surely aware of *a* correct explanation of the rope-snapping. If know which laws were involved, then I surely have a much better explanation of that same event: I have a *scientific* explanation of that event, whereas I previously did not.

Causal knowledge as differential knowledge

Earlier it was pointed out that we seem to know of individual cases of metal's being heated and expanding for that reason *before* we know of the law in virtue of which heated metals expand: indeed, a precondition for our knowing of such a law seems to be that we know of individual cases of metal expanding in consequence of being heated.

In response to this, the following is likely to be said (and---though not as a response to this very point---*has* been said, innumerably many times, by advocates of DN, from Mill to Ernest Nagel):

> *(#) We see many cases of metal objects being heated and then expanding. We are puzzled. So we posit a law L to the effect that metal expands when heated. Now that we (think we) know of this law, we are no longer puzzled.*
>
> *You say that we can't know of L's existence until we know of individual cases of metal objects' expanding in consequence of being heated. This is not correct. We know of metal objects' expanding upon heated. But until we know of L, we don't know that there to be a cause-effect relationship between those objects' being heated and their expanding.*

L is presumably to the effect that, for any metal object x, x expands if heated. If L doesn't support counterfactuals, it isn't a law. So L must support counterfactuals. But if it supports counterfactuals, L is a causal claim (*for any metal object x, heating x causes it to expand*). So let us suppose, as we must, that L is a causal claim.

If our only reason for thinking that x expands when heated is that x *does* expand when heated, then in positing a law to the effect that they are causally connected, all we are doing is saying: "It is so; for it must be so."

In which case, L doesn't really explain anything. And in which case, furthermore, there is insufficient evidence for L, a consequence being that L isn't really a law so much as it is a tentative hypothesis, which, being of questionable veridicality, cannot serve as the basis of an explanation.

By the same token, if we have good reason to believe L to be a law, we must have reason to believe, independently of L, that this particular piece of metal expanded because it was heated, the same being true of that piece of metal piece, and also of that other piece of metal, and so on. How such knowledge is acquired is made clear over the course of the next ten sections.

Causation as proportionality restoration

Laws are usually expressed by dependence-relations:

**The size-increase depends on the temperature increase.*

**The pressure-increase depends on the volume-decrease.*

**The one object's acceleration depends on its mass and on the momentum of the object that collided with it.*

**The increase in x's mass, relative to a framework with respect to which x is accelerating, depends on the rate of acceleration.*

**The extent to which electron x's position can be known depends inversely on the extent to which its velocity can be known, and the extent to which x's momentum can be known depends inversely on the extent to which its position can be known.*

**The extent to which exercise lowers blood pressure depends on the amount of exercise done.*

**The extent to which insults wound a person depends on the extent to which he gives them credence.*

**The extent to which poverty leads to crime depends on the extent to which poverty leads to food-shortages.*

**The extent to which a person needs eight hours of sleep on a given night depends on the extent to which he is running a sleep deficit.*

**If a solid object is rotated around one of its end points, the extent to which it remains straight depends inversely on the length of the object.*

A related point is that if a given force (e.g. gravity) is acting on an object, it doesn't follow that the object's state of motion changes or that the object otherwise changes. If another force, e.g. an opposite and equally strong gravitational force, is acting on it, the resulting movement on the part of the object will be the vector sum of both forces. (Whence the distinction, notional if not actual, between kinematics and dynamics.) The object may therefore remain stationary or it may move, in any direction with any degree of rapidity. [8]

Physical laws as regularities of proportion

Threshold effects aside, effects are proportional to their causes. This has several consequences, none of them consistent with RA or with DN.

First of all, laws are not expressed by regularities. A collision may or may not displace an object. It may be destroy it. It may not destroy it, but may fail to make it budge. If it does displace it, it could move, with any degree of rapidity, along any one of a number of a possible trajectories. If conditions are held constant, then events that are identical in respect of momentum or temperature (or whatever the relevant parameters are) have the same effects, in which case causal relations are expressed by regularities. If there is to be a perfect regularity of some kind, there must an event e such that perfect e-duplicates occur over and over. That condition isn't met if conditions are at all variable, since inputs vary when conditions vary.

Second, to the extent that laws *are* expressed by regularities, they are expressed by regularities of *proportion*: volume increases directly proportional to temperature increases; gravitational attraction is proportional to the masses of the bodies in question; if an object is subject to no differential forces, the degree to which its trajectory deviates from a straight line depends on the length of that trajectory.

So if indeed causal nexuses are regularities, they aren't regularities of the form *e-similar events always precede e*-similar events.* They are regularities of proportion and are therefore regularities of a much more abstract kind than those that the DN-advocate has in mind.

Explanation of the fact that statistics tend to track causal substructure

In some cases, we may get lucky: conditions may remain perfectly constant; a series of events occurs each member of which is qualitatively and quantitatively just like each of the others. In that case, the effect of any given one of those events will be the same, qualitatively and quantitatively, as the effect of any other given one of those events.

But circumstances don't remain constant, and they don't have to remain constant for us to learn what causes what. Events are changes. The greater the change, the greater then effect. As long as we can correlate input-differences with output-differences, we can establish causality.

A corollary is that a singular events are the effects of a singular antecedent, and, when given a singular event, we know its cause as soon as we identify a singularity among its antecedents.

Example: You are a square: somebody who follows all the rules. One night, after a lifetime of celibacy, you decide to sleep with a lady of the night. A few days later, you have a venereal disease.

Singular event: the malady. Singular antecedent: the dalliance. Effect: disease. Probable cause: The dalliance.

Do such dalliances always lead to venereal disease, or even most of the time? No. Does it matter? No. You know of some event predating your malady that is as singular as your malady. Were that event not the cause of your malady, you would have a genuine anomaly on your hands; but since it is, you don't.

Of course, if you had dalliances with 35 such women, then of course you would not be able to pinpoint the cause---but only because any given one of those events could be the cause, notwithstanding that not a single one of them was an instance of an event-type whose instances are always, or even usually, followed by the onset of disease.

And this proves the point at hand, that, when e is known to be the cause of e*, it isn't because e-e* is an instance of a perfect regularity; it's because, of e*'s various antecedents, e is the only one that is proportional to e*.

If e* has many antecedents $e_1 \ldots e_n$, each of them proportional to e*, then e* has many possible causes. In that case, the actual cause has to be decided experimentally. For each i ($1 \leq i \leq n$), does an e_i-similar lead to an e*-similar event, under conditions similar to those under which e brought about e*? Supposing that the answer is in each case "no," then, pending further information, e must be presumed to be the cause of e*.

Another example: You force yourself to run three miles. You have a heart-attack.

Singular event: the heart-attack. Singular antecedent: you run three miles. (You've never run so much as a 100 yards in your life.)

Do all or even most physically unfit people who run three miles have heart-attacks? No. But there needn't be such a regularity for the jog to have brought on the heart-attack, and you don't need to know of such a regularity, to know what brought on your heart-attack.

Given that we can know what causes a given event without knowing of any law L that links cause and effect, it follows, given that to explain an event is to identify its cause, that we can sometimes explain events without explaining them in terms of laws.

To explain an event *in detail* is to explain how the specifics of the event depended on the specifics of the cause, and for that we *do* need laws. We need laws if we are to have *scientific* explanations, Where SE is a scientific explanation of an event e just in case:

(i) SE identifies the event e that caused e,*
(ii) SE identifies the law in virtue of which e caused e*; and, so it follows,*
(iii) SE identifies just what it was about e that was responsible for its bringing about e.*

But scientific explanations presuppose prescientific explanations. We must distinguish between non-explanations and pre-scientific explanations. Non-explanations are not pre-scientific explanations. If my explanation of the fact that I burnt my hand is that I touched a hot surface, my explanation is correct, but not scientific. If my explanation is that my parents weren't home and it wasn't my touching the hot-stove *per se* but my transgressing their instructions (*don't touch the stove*) that caused me to feel scorching pain, then my explanation is false. Also, if my parents give me permission to touch the stove (thinking that it's off), and (because it's on) I scorch my hand, I'll know that, indeed, it was, on each occasion, because the stove was hot that I burnt on my hand; and on neither occasion is it necessary that I know of some regularity to the effect that all people who touch hot stoves feel pain.[9]

The role of experimentation in the acquisition of causal knowledge

If it is always so obvious what causes what, why do scientists bother with tests? Why do they repeat the same experiment over and over?

Scientific tests have a number of functions. First of all, people's causal beliefs aren't always correct. This is partly due to the fact that, as was pointed out in class, a given event may have a number of equally salient antecedents. To a larger extent, however, it's due to the fact that, when people look for causes, they have some tendency to look for causes that fit a pre-existing narrative and, to that extent, tend not to look for the actual causes. "I didn't get the job: my ex-boss must have blacklisted me." I'm overlooking the fact that I was 45 minutes late to the interview.

Another reason is that experiments identify specific information about causes that is not easy to come by in non-experimental settings. It's easy to know that a monkey-bite made you sick. It's hard to know exactly what it was that the monkey transmitted to you. And it won't be known if experiments aren't done (not necessarily on you personally, but on other monkey-bite victims).

Also it would be hard to know under whether the pathogen that was transmitted to you would, if transplanted into another person's bloodstream, affect that person: maybe your immune-system is unusual in some way.

Further, it would be hard to know whether or not a bite, or some other way of intravenously introducing the pathogen, is the only way of doing so: Is the pathogen airborne? If not, is it transmitted through casual contact? These questions cannot be answered except on the basis of controlled studies. At the same time, those controlled studies presuppose correct if approximate beliefs as to what, in some given case, led to some malady relevantly similar to yours.

Further, supposing that events of kind cause events of some other kind, it's only in controlled environments that the exact nature of the operative dependence-relation can be learned. It's easy to figure out that hitting a ball makes it move. It's hard to figure out the specifics. It's hard to identify a true, general principle to the effect that, given an object x, if x has such and such a structure and can therefore withstand impacts of a certain amount of force without being deformed, then, neglecting friction and air-resistance, if x is struck by an object with momentum m, then x will be displaced by a certain amount.

The instantaneousness of the interval between cause and effect

If e is the cause of e*, e precedes e*. Given that no event is instantaneous, this means that each takes time to occur and, therefore, that it takes time for both of them to occur.

Let INT be the interval during which e and e* happen. Events could have occurred after e that prevented e* from occurring or at least prevented it from occurring in the exact way it did. Ball A collides with ball B. The collision causes B to move to the edge of the table. During the time that it took B to travel to the table-edge, events could have intervened (somebody could have grabbed B, it might have been subjected to extreme heat and disintegrated as a result, and so on) that would have prevented B from doing so. Supposing that the distance from the collision to the table's edge is n meters, given any spatial interval less than n, no matter how small, events could have occurred that would have prevented B from traveling that distance or from traveling that distance along that particular trajectory.

Thus, given any non-null time-interval, interval-initial circumstances of a given kind are compatible with interval-terminal circumstances of many different kinds. The shorter the intervals one considers, the greater the guarantee the probability that interval-initial events will have these as opposed to those interval-terminal outcomes. But that probability is always less than 100%. It approaches 100%, as one considers increasingly small intervals. But, since nothing happens during null time-intervals, the probability that a given event will eventuate in this as opposed to that sort of event is always less than 100%.

Therefore, instances of causation are not instances of regularity. For e to be the cause of e* is not for e-similar events to always precede e*-similar events. Whatever it is to be an "e-similar" or an "e*-similar" event, e-similar events---though causative, in some contexts, of e*-similar events---are not necessarily always followed by e*-similar events.

This in turn means that, if L is the law in virtue of which e caused e*, L is not to the effect that e-similar events always precede e*-similar events; indeed, L is not to the effect that events of *any* kind follow events of some other kind. It follows that laws are not regularities. It also follows that laws are not "forces" or "structures" that are responsible for regularities.

Natural Laws dispositive of, but not identical with, regularities

The law of inertia *disposes* bodies have constant velocities during arbitrarily long intervals. And, given some body x, so long as x's state of motion is governed by no law other than the law of inertia, its state of motion will indeed remain constant. Further, if x is acted upon by many forces other than the force of inertia[10], x's state of motion will *still* remain constant *so long as* those other forces cancel one another out.

It was on the basis of his knowledge of this fact that Galileo discovered the law of inertia. Galileo considered bodies that, apart from their own inertias, were in states of varying degrees of equilibrium; and he found that that, the closer the equilibrium in question was to perfect equilibrium, the more the projectile in question was to have a constant state of motion (during the time-interval during which that equilibrium lasted).

It was on the basis of this same principle that the Inverse Square-law (*the gravitational attraction between x and y is directly proportional to their masses and inversely proportional to the square of the distance between them*) was definitively verified. An environment was constructed in which each of two balls were in a state of perfect equilibrium, except with respect to the gravitational attraction of the other ball: and the balls moved in exactly the way that, according to ISL, they would move if no forces, other than the gravitational pulls they exerted on each other, were acting on them.

But no physical object is isolated from all others. And if a given physical object is in equilibrium with respect to all forces but one, it is seldom for very long. For these two reasons, it is only under exceptional circumstances that laws give rise to regularities.

How we know what causes what

According to Hume, we know that e causes e* because

(i) e immediately precedes, and is adjacent with, e,*

and

(ii) e-similar events always immediately precede adjacent e-similar events.*

This is not typically how we learn what causes what. If it is known that, apart from being collided with by ball A, hitherto stationary ball B was in a state of equilibrium, then it is known that the collision was necessarily the cause the change in B's state of motion.

"But can it ever be known that a given thing is in equilibrium in every respect but one?"

It can be known that, in every respect but one, a given thing is in a state that is sufficiently equilibrium-like that there was but one possible cause of some change subsequently undergone by that thing.

If it is known that, apart from the nasty remark I just made about your cooking, you were in a state of relative psychological equilibrium, then, even though most people whose cooking is insult don't go into a fit of rage, I know that *your* outburst was indeed the effect of my remark.

Thus, if we know that e is the cause of e*, it isn't necessarily because we have observed some regularity of which that event-sequence (e followed by e*) is an instance. In many cases, there is no such regularity. Cf. the just-discussed tantrum.

And in many cases where there is such a regularity, it isn't on the basis of our knowledge of that regularity that we know e to be the cause of e* Cf. the situation.

A child touches a hot stove and feels excruciating pain. He knows on that basis alone that touching a hot stove will lead to excruciating pain. He knows this because he knows that, apart from the fact that he was touching the stove, there was nothing distinctive about his situation. It is true that, whenever people touch hot stoves, they feel excruciating pain. So there is in this case a regularity of which the event-sequence in question is an instance. But it isn't on the basis of the child's knowledge of such a regularity that he knows what caused his pain.

Knowledge-why=high-resolution knowledge-what

It may be that, unless one has knowledge of laws, one cannot know *why* a given event followed some other event. But without knowledge of laws one can obviously know *that* the first event was indeed the cause of the second.

Hume seems to have thought that, unless it is known *why* the collision caused the hitherto stationary ball to move---unless Newton's second law of motion is known, along with A's momentum at the time of the collision and B's mass---- it cannot be known *that* the collision is the cause of its moving; that in order to know *that* e is the cause of e*, one must know of some 'covering law' of which that event-sequence is an instance. And this, we have seen, is false.

Singular causation and the falsity of the Deductive-Nomological Model of explanation

To know what caused an event is to know how to explain it. I can explain your tantrum: I insulted you. There are certainly better explanations ---one that might be known to your psychoanalyst: but according to each of those explanations, the insult triggered the tantrum.

The difference between my explanation (*I insulted you*) and Freud's *(given your desire to fill the shoes of your father, who is a revered chef, and given your high opinion of my cake-related opinions, the insult initiated a bout of self-loathing, reinforced by loathing of your father for imposing such high standards on you, whose intensity was so extreme that no non-tantrum-throwing behavior could have discharged it adequately*), is that the latter is more high-resolution than the former. But the latter includes the former -- just as Galileo's explanation of the change in B's state of motion will include the layperson's explanation.

This shows that, contrary to what DN assumes, not all explanations are law-based. In some cases we know that e causes e*, and we are therefore know a (low-resolution but otherwise) correct explanation of the latter, even though we don't know of any law that connects the two events.

It also shows that knowledge of (low-level but otherwise) correct explanations of facts is a prerequisite to knowledge of natural law. In order to learn Newton's second law (*the acceleration of an object depends directly upon the net force acting upon that object and indirectly on the object's mass*), one must know that certain collisions are responsible for certain displacements. If one doesn't know this, then one *a fortiori* doesn't know that among the (invariably very numerous) laws determinative of the exact nature of the accelerations of the displaced objects, one of them is that the acceleration of the displaced object is inversely proportional to its mass and directly proportional the net force acting on it.

The failure of the Deductive-Nomological model in connection with psychological explanation

One quite glaring failure of DN concerns its treatment of psychological explanations. If DN is right, as Hempel makes clear, not a single one of the psychological explanations ever put forth identifies the cause of any psychological event. If DN is right, I don't know that my insult is the cause of your tantrum unless I know of some perfect regularity of which that event-sequence is an instance. I know of no such regularity. Outside of physics and chemistry, any such regularities as are to be found are imperfect. And a consequence of DN, therefore, is that, outside of those two disciplines, nothing has ever been explained.

This must be understood aright. It may well be that nobody knows the complete explanation of any given psychological event. But a consequence of DN is that, pending the discovery of some perfect regularity of which the insult-tantrum event-sequence was an instance, it simply isn't known whether or not the insult was in any way, to any degree, the cause of the tantrum.

And this consequence of DN is simply false. The cause of many a psychological event is known, even though the corresponding regularity, if there is one, is not known; and if the corresponding regularity is known, it isn't necessarily implicated in one's knowledge of the relevant event-sequence. You may well know that people who are dumped feel sadness. But it isn't on the basis of this knowledge that you know that the reason you are sad is that you were dumped. And you know that you intend to write a five-voice fugue in order to impress your ex-girlfriend, so as to rekindle her interest in you. But even though dumped people often (though not always, or even usually) do things to impress the people who dumped them, they seldom, if ever, try to do so by writing five-voiced fugues.

DN entails that, unless you know of some regularity of which your writing this fugue is an instance, you cannot possibly know that your doing so is a consequence of your being dumped. DN also entails that, unless people who are dumped always respond by trying to impress their exes, no person's attempting to rekindle his/her ex's affections through a heroic deed is a consequence of him/her being dumped by said ex---which is false---and *a fortiori* it cannot be known of some particular act of valor performed by some dumped person that his/her heroic deed is a consequence of him/her being dumped---also false.

DN in relation to statistical explanation

Hempel acknowledges that perfect regularities often are not available. And he says that, in such cases, *one* must rely on imperfect regularities: regularities of the form *e-similar events usually precede e*-similar events.*

This move is misconceived. The probability that an insult will elicit a tantrum is small: most insults do not result in tantrums. And supposing that 95% of insults were followed by tantrums, there would not, at least not for that reason, be a 'statistical law' to the effect that 95% of tantrums follow insults; for it might not be that the tantrums in question were consequences of the antecedent insults.

In general, supposing that most phi's are followed by psi's, it doesn't follow that phi's cause psi's, even part of the time, or *a fortiori* that there is a statistical law to the effect that n% of phi's cause (or are followed by) phi's (for some n>50).

For example: most people who are swimming pools wear some sort of bathing gear (i.e. they don't swim in the nude). But that isn't because they're in a swimming pool; it's because (in most cases) they're in a public place. They are often, presumably, wearing bathing gear *despite* being in a swimming pool.

Contrariwise, given only that 30% of phi's are followed by psi's, it doesn't follow that phi's *don't* cause psi's. Smoking causes emphysema; but only emphysema occurs in less than 30% of people who smoke the requisite amount (the amount the emphysemics smoked, at least before they got emphysema).

What *is* true of emphysemics is that the only *distinctive* thing about their behavior was that they were heavy smokers. For this reason, Wesley Salmon introduced the 'statistical relevance model,' as an alternative to DN. But the statistical relevance model, being nothing more than a watered down version of the Deductive nomological model, has the latter's defects but not its virtues; and it's not really worth running the diagnostics on this obviously degenerate and degenerating research program.[11]

Delimiting DN's sphere of application

DN doesn't hold with respect to all explanations; it isn't possible to give deductive nomological explanations of events until it is known what is causally relevant; and to the extent that one knows what is causally relevant to what, one already has knowledge explanations.

So far as DN is correct, it's a correct description of *scientific* explanation, a "scientific explanation" of E being one that says both

(i) What *made E happen*

and

(ii) How *it did so.*

[12] The Raven-Paradox

As a matter of logic, sentences that are logically equivalent must be confirmationally equivalent, but there are logically equivalent sentences that are not confirmationally equivalent (e.g. 'all ravens are black' and 'all non-black things are non-ravens').

Logical equivalence: S1 and S2 are logically equivalent iff S1 entails S2 and S2 entails S1. (n=4 entails n=3+1, and n=3+1 entails n=4. Therefore, n=4 is equivalent with n=3+1.)

Confirmational equivalence: S1 and S2 are confirmationally equivalent iff, given any datum D, D confirms S1 to degree n iff D confirms S2 to degree n. (D: Two days ago, Smith had no cars, but yesterday he won four cars on a game show. D confirms n=4 and n=3+1, and it confirms them equally.)

Logical and confirmational equivalence logically equivalent: For S1 and S2 to be logically equivalent is, by definition, for neither to entail anything not also entailed by the other. It follows, by an indirect proof, that logical equivalence entails confirmational equivalence.

Suppose that S1 and S2 are logically equivalent but confirmationally non-equivalent. Since they are confirmationally non-equivalent, there is some proposition S3 that confirms S2 more than it confirms S1. It follows that S3 is relatively unlikely to be false if S2 is true and also that S3 is relatively unlikely to be true if S1 is true. Thus, S2 entails that S3 is relatively likely to be true and S1 entails that S3 is relatively likely to be false. Thus, there is some proposition that S1 entails that S2 does not entail and there is some proposition that S2 entails that S1 does not entail. Thus, S1 and S2 are not logically equivalent, contradicting our hypothesis. Thus, confirmational non-equivalence entails logical non-equivalence. Equivalently, logical equivalence guarantees confirmational equivalence.

The Raven Paradox: Consider:

R1: All ravens are black (if x is a raven, then x is black), and
R2: All non-black things are non-ravens (if x is non-black, then x isn't a raven).

R1 and R2 are equivalent.

(That is, they are equivalent if read *extensionally*, i.e. as concerning only actual, as opposed to

possible objects or, equivalently, as concerning what *is* as opposed to what *must be*. Hempel doesn't acknowledge that R1 and R2 can be read either extensionally or intensionally. When expounding Hempel's theory, we will follow Hempel's example in this regard. When critiquing it, we won't.)

The (extensional) equivalence of R1 and R2 is easily verified using Venn diagrams. But there is a problem. Consider:

R3: x is a pink piano.

R3 confirms R2, but not R1. The pinkness of pianos has nothing to do with the blackness of ravens. But R3 would confirm R1 and R2 equally if they were logically equivalent. Since it doesn't, they aren't. But they seem to be ---a paradox.

Another paradox of the same kind: Consider:

M1: Heated metal expands (if x is metal, then x expands when heated), and
M2: Heated non-expanders aren't metal (if x doesn't expand when heated, then x isn't metal).

The equivalence of M1 and M2 is easily verified using Venn diagrams. But there's a problem. Consider:

M3: Cube steaks shrink when heated.

M3 confirms M2, but not M1: The deflationary tendencies of heated meat have nothing to do with the expansionary tendencies of heated metal. But M3 would confirm M1 and M2 equally if they were logically equivalent. Since it doesn't, they aren't. But they seem to be a paradox.

Hempel's solution

R3 *does* confirm R1 and R2 equally. We think otherwise because, our background beliefs being what they are, we believe it impossible that pianos should be ravens. We wouldn't believe this if our background beliefs were different, and we'd therefore see that R3 *per se* is equally confirmatory of each of R1 and R2.

Similarly, M3 *does* confirm M1 and M2 equally. We think otherwise because, our background beliefs being what they are, we believe it impossible that cube steaks should be made of metal. We wouldn't believe this if our background beliefs were different, and we'd therefore see that M3 *per se* is equally confirmatory of each of M1 and M2.

The problem with Hempel's solution

We can't learn anything about ravens by studying pianos. We can't learn anything about metal by studying meat. Hempel's solution implies otherwise and is therefore false.

The actual solution to the Raven Paradox

There are two ways to interpret:

(1) All phi's are psi's (if x is a phi, then x is a psi).

(1) can be interpreted (or 'read') *extensionally* or *intensionally.* Read intensionally, 1 means:

(1) For reasons of natural law, all phi's are psi's.*

Read extensionally, 1 means:

(1#) By coincidence, all phi's are psi's.

Some statements, for example,

(2) All of the coins in JM's pocket are quarters

are naturally read extensionally, e.g. 2 is naturally to be taken to mean:

(2#) By coincidence, all of the coins in JM's pocket are quarter.

But if there were some natural law (or mechanism---a mechanism being a way deploying natural laws in furtherance of some objective) that required coins in my pocket to be quarters, then (2) would have to be read intensionally and thus taken to mean:

(2) For reasons of natural law, all the coins in JM's pocket are quarters.*

M1 is naturally read intensionally and is thus taken to mean:

(M1) For reasons of natural law, metal objects expand when heated.*

Read extensionally, M means:

(M1#) By coincidence, all things that don't expand when heated aren't metal.

M1* is *not* equivalent with M2. M2 is to the effect that a thing's failing to expand when heated implies that it *already* isn't metal, and it therefore implies that its failing to expand when heated cannot possibly *cause* it not to be metal.

Thus, M1* entails the negation of M2 and *a fortiori* isn't equivalent with M2.

An analogous argument establishes the non-equivalence of R1 and R2.

The Raven-paradox a by-product of the Regularity Analysis

For Hempel, as for Hume, there are no causal connections, only regularities that mimic causal connections. So, for Hempel/Hume, any connection between

x is heated metal

and

x is expanding

is purely extensional.

If that were indeed so, then a shrinking cube-steak would provide quite as much support for *x is heated metal, then x is expanding* as would a piece of heated, expanding metal.

But it doesn't; so there is no paradox.

Chapter 6 The Problem of Induction

What is the problem of induction?

P: The sun has always risen.
Q: It will rise tomorrow.
U (The Uniformity Principle)*: Nature is uniform. More precisely: In its more general respects, nature doesn't change. Even more precisely: Although there are changes in the specific way in which mass-energy is distributed, there are no changes in the laws governing such changes.*

P doesn't entail Q. So it is only to the extent that P *confirms* Q that P provides any support for Q. And it is only to the extent that U is the true that P confirms Q. We know that nature *has* been uniform. But it's only if nature's past behavior is representative of its future behavior that nature *will* be uniform, and it's only if nature *will* be uniform that U is true. If we assume that nature will be uniform, given that thus far it has been uniform, the inference we are making is an exact analogue of the inference from P to Q. In fact, the inference from P to Q is a special case of that inference. Since the very thing we are trying to show is that P confirms Q, we can't assume that nature's past behavior is any indication of its future behavior; for the assumption that P confirms Q is implicit in that assumption and therefore provides no independent support for it.

The problem with the problem of induction

We must distinguish *induction by enumeration* (also known as *enumerative induction* (EI)) from *inference to the best explanation* (IBE).

Enumerative induction (EI) *defined*: Every known phi has also been a psi. Therefore, unknown phi's are also psi's. (Every known raven has been black. Therefore, unknown ravens---those we will come to know of but don't of yet, and those we will never know of---are black.)

Inference to the best explanation (IBE) *defined*: Let $d_1 \ldots d_n$ be the data at our disposal. Suppose that, if $d_1 \ldots d_n$ were all there was to know about the world, then the world would contain many anomalies, i.e. many unexplained phenomena. Let H be the supposition that X exists, where X is something not yet known to exist.[13] Supposing that H is put forth legitimately, H is put forth because, if correct, it eliminates some of the aforementioned anomalies. H is thus a case of IBE.

In general, an IBE is a hypothesis to the effect that some hitherto unknown entity X exists such that, if X does indeed exist, then some hitherto anomalous phenomenon ceases to be anomalous.

Example: Smith has trouble reading certain words, even though he is otherwise highly intelligent. Hypothesis: Smith has dyslexia. If that hypothesis is correct, then it is no longer anomalous that he has trouble reading those words, even though he is a person of a high intelligence.[14]

No such thing as enumerative induction

In some cases, a person's acceptance of an IBE, coupled with knowledge on his part that known phi's have also been psi's, disposes that person---sometimes legitimately, sometimes illegitimately---to hold that (as of yet) unknown phi's are also psi's.

However,

(1) Unless coupled with acceptance of an IBE, one's knowledge that past phi's have been psi's doesn't legitimate his believing that unknown phi's are also psi's.

Indeed,

(2) Unless supplemented by knowledge of an IBE, one's knowledge that past phi's have been psi's does not lead to his believing that unknown phi's are psi's. Whenever someone seems to infer

(a) All phi's are psi's

from

(b) All known phi's are psi's,

that person is really inferring (a) from (b) coupled with some IBE.

(1) is an epistemological (or logical) claim.
(2) is a psychological claim.
Each is demonstrably correct.

(1) Defended

Suppose that, for every x, such that x is an object known to be metal that is heated, x is also known to expand upon being heated. (I.e. suppose that a purely extensional (*de facto*) as opposed to intensional (principled, causation-based, *de jure*) connection has been observed to the effect that *if x is metal and x is heated at t, then expands at t*, where t* follows t by a vanishingly short interval.*) Further, suppose that, on the basis of this fact alone, it is hypothesized that, if unknown metal objects are heated, they too will expand. Such a hypothesis would be plainly spurious. If it is suspected that heating a metal object *causes* it expand, *then* H is reasonable. If it is suspected that, although heating a metal object doesn't cause it to expand, nonetheless the thing that causes metal to be heated also causes metal to expand, *then* H is reasonable. If it is believed that a metal object's being heated isn't causally related to its expanding, then it is *ipso facto* believed to be nothing more than coincidence that known metal objects expand when heated, and it is therefore not reasonable to believe that unknown metal objects expand when heated.

Another illustration of this principle: I roll a certain die 10 times in a row. Each time I roll a six. Here I can choose between two hypotheses.

> *H_1: It was pure coincidence that each roll was a six roll. The die was fair.*
> *H_2: There is a causal mechanism that disposes the die to land six-side-up. (The die is weighted. The side opposite the six-side is magnetically attracted to the surface on which I'm roll the die. Etc.)*

If I accept H_1 and on that basis alone hypothesis that future rolls will be six-rolls, my hypothesis is irrational, as it is given by the proposition: *There is nothing that makes a six-roll more probable than an n-roll ($1 \leq n \leq 5$). Therefore, the likelihood that the next roll is a six-roll does not exceed the probability that it is an n-roll. Therefore, the likelihood that the next roll is a six-roll* does *exceed the probability that it is an n-roll.*

If I accept H_2 and on that basis hypothesize that the next roll will be a six-roll, my hypothesis is rational, as it is given by the proposition: The laws of nature being what they are, and the operative conditions being what *they* are, the next roll's not being a six-roll would violate the laws of nature.

(2) Defended

Smith has been rolling sixes all night, and Smith believes that he will continue to roll sixes. Smith is not a man of science. He doesn't think in terms of laws of nature and causal mechanisms. So why does Smith believe that his lucky streak---if that's what it is---will continue? Precisely because, in his opinion, it *is* a lucky streak: Smith believes in Lady Luck. He doesn't know anything about physics. But he does believe that the Gods favor certain people at certain times; and he believes---rationally, relative to the position that the Gods exist---that the Gods call the shot. So Smith's belief in the continuation of his lucky streak is not a case of enumerative induction; it is an IBE, namely: *Given that I've been so lucky, Lady Luck has my back tonight, and She will therefore make sure that I continue to roll sixes.*

Question: What if Smith doesn't have any reason at all for believing that he'll continue to roll sixes? What if his six-roll expectations are a consequence of Pavlovian conditioning: He has come to associate die-rolls with six-rolls; this psychological association causes him to think *six-roll* whenever he thinks *die-roll*, which in turn causes him to think *I will roll a six* whenever he thinks *I will roll the die.*

Answer: In that case, Smith's confidence that his future die-rolls will be six-rolls doesn't embody any inference of any kind. Not enumerative inductive, not inference to the best explanation. Human nature being what it is, Smith is likely to produce a *post hoc* rationalization of this expectation of his: he's likely to believe, at some level, that Lady Luck is on his side, or some such. But, so far as his expectations concerning future rolls are really the result of Pavlovian conditioning, they don't embody inferences of any kind. Inferences are judgments. Associations are not judgments---though people tend to rationalize their existing associative tendencies by superimposing judgments on them.

Question: What if Smith believes wrongly that enumerative induction is a viable form of inference? What if he believes that, even if there is no causal mechanism disposing a given die-roll to be a six-roll, he believes that, since his past rolls have been six-rolls, they must, for that reason alone, continue to be six-rolls?

Answer: As previously explained, if that is what Smith believes, then Smith believes that he *isn't* more likely to be roll sixes than he is to roll n's ($1 \leq 5 \leq n$). If, in addition to having *that* belief, Smith also believes that he *is* more likely to roll sixes than he is to roll n's, then the part of his mind that causes him to roll the die isn't the part of him that believes that he isn't more likely to roll sixes than he is to roll n's: the operative part of him *does* believe that he's more likely to roll sixes than he is to roll n's. And to the extent that his past six-rolls are dispositive of future six-rolls, he *ipso facto* believes in some mechanism connecting his past rolls to his future rolls; and his confidence that he'll continue to roll sixes is therefore IBE-based.

EI is unjustifiable. (Hume said why.) When people seem to make EI-based inferences, they are not in fact doing so. In some cases, their mentation is reflexive and thus doesn't embody inferences of any kind. In others, they have two beliefs: one to the effect that the sole relevant fact is the fact that , although all known phi's have been psi's, there is no non-fortuitous connection between a thing's being a phi and its being a psi and there is therefore no reason to believe that future phi's will be psi's; the other to the other effect that, since all known phi's have been psi's, there is some mechanism that disposes phi's to be psi's---it could be the Gods, it could be the laws of physics, it could be (and often is, especially where gamblers are concerned) an infantile belief in the 'omnipotence of thought.'[15]

Discontinuity-elimination the essence of explanation

If correct, a given IBE is a case of anomaly-elimination: supposing that $d_1 \ldots d_n$ are the empirical data at our disposal, H is a good IBE to the extent that H fills in the gaps left by $d_1 \ldots d_n$.

Question: What if $d_1 \ldots d_n$ are such that any hypothesis H that fills in the gaps left by *them* is inevitably false? What if $d_1 \ldots d_n$ is such a limited or otherwise imperfect data-set that, if a hypothesis does a good job of modeling that particular data-set, it stands no chance of being true?

Answer: This happens. But we have to work with the data we have, not the data we want. (*The data we have*: what our eyes, ears, etc. have told us. *The data we want*: the data that, if appropriately modeled, apprises us of the structure of the world.)

Question: How do we deal with this fact?

Answer: By acquiring more data. False hypotheses may do a good job of modeling limited data-sets. But, precisely because they're false, they tend not to model enlargements of such sets. False hypotheses don't travel well. A lie may work in particular context: you lie to your roommate about eating his powdered doughnuts. But as data comes in, it becomes increasingly difficult to compatibilize the lie with the facts: he turns on the light and sees doughnut powder on your face, the doughnut jelly on your fingertips, etc. The jig is up.

Question: Given some relatively large and comprehensive data-set $d_1 \ldots d_n$ ---some data-set that is sufficiently large that false hypotheses are unlikely to model it---what is it to 'fill in the gaps' in that data-set?

Answer: To eliminate anomalies is to eliminate discontinuities. To explain an event e* is to identify some e that is the cause of e*. e is the cause of e* if, given e, e*'s occurrence is not discontinuous with the state of the universe prior to e*'s occurrence.

Example: You see a red ball roll into an opaque (non-see-through) tunnel. A moment later you see a red ball (that is otherwise the same as the first ball) roll out the other end of the tunnel. Hence your data-set:

d_1: Red ball rolls into one end of tunnel, located at place p, at time t.
d_2: Red ball rolls out of other end of tunnel, located at place p (10 feet away from from p), at time t* (10 seconds after t).*
d_3: What you saw, or otherwise perceived, during the spatiotemporal interval defined by the distance between p/t and p/t*, is in no observable way continuous with the sudden emergence, at p*/t*, of a red ball.*

Hypothesis (H): The ball that rolled into the tunnel is identical with the ball that rolled out of the tunnel; and while the ball was in the tunnel, it continued to roll, that being why it was able to exit the tunnel, even though it was not observed to roll.

If the totality of the facts were as d_1-d_3 described them, the emergence of a red ball from the tunnel's end would be an instance of *ex nihilo* creation: an event that had no roots in the state of the universe prior to its occurrence. Given both d_1-d_3 *and* H, the ball's exiting the tunnel is appropriately continuous with the state of the universe prior to its occurrence.

The reason is that, given d_1-d_3+H, the ball's emerging at t* from the tunnel is a part of a continuous process.

Continuity=causality.

Explanation=continuity-restoration.

To model data $d_1 \ldots d_n$ is to posit something X such that, given d_1-d_3+X, the facts described by $d_1 \ldots d_n$ are fragments of structures/processes that don't comprise discontinuities.

Inductive knowledge as knowledge of continuities of continuities

Question: What right have we to suppose that the world is continuous? What right have we to posit invisible continuities underlying the visible discontinuities of which observation apprises us?

Some non-answers: Let us first talk about how *not* to answer this question.

We cannot say

> (H#): In the past, the unobserved continuities that we posited later came to be observed; so, inductively speaking, it is a good idea to posit continuities.

H# is off-limits, for the very reasons Hume gives. Nothing we say can presuppose that the past is a good indicator of the future; otherwise, Hume's Argument (HA) will successfully nullify our proposals.

Further, we cannot say that

> *(H%): There are laws of nature; we know what those laws are; and those laws being what they are, the future has to be like the past, at least in such respects as enables our knowledge of the past to support certain cases of inductive inference.*

HA nullifies H%. Suppose *arguendo* that there are laws of nature that, given the present state of the universe, fix the future states of the universe. Given HA, we don't know what those laws are unless the data-sets on the basis of which we put forth our theories as to the identities of those laws is similar to the data-sets that we will have in the future. We can't know that past data-sets resemble future data-sets unless we know of laws that guarantee that they'll resemble each other; and we can't know of such laws unless we know that they'll resemble each other.

The justification for Inference to the Best Explanation

Remember that IBE is continuity-restoration. Bearing this in mind, suppose that, at time t*, nothing that occurs is in any way continuous with what occurred at any prior time t. Let U1 be the pre-t* universe and let U2 be the post-t universe. On what basis could it be said that U1 and U2 were both the *same* universe, that there was some one space-time manifold U such that U1 and U2 are earlier and later parts of U? By supposition, nothing in U2 would be a continuation of anything in U1. None of the memories (or mental contents that, if veridical, would be memories) would be continuations of anything that happened in U1.

People might have memory-like mental contents to the effect that, during the pre-t* period, such and such happened. But those mental contents wouldn't really be memories. By hypothesis, they would have no roots in U1; there would be nothing that happened in U1 that, at t*, wasn't cut off; and there would thus be nothing that happened in U1 that left a record R in some pre-t* person's mind, such that R continued to exist until t* and therefore continued to persist. By hypothesis, R was cut off at t*. Any post-t* R-similar mental content wouldn't be R, just as, if your car were destroyed right now, and a similar car were built later, that new car wouldn't be identical with your old car, owing the fact that the new car wasn't a continuation of the old one.

For the same reason, no non-psychological post-t* state of affairs would be a continuation of any pre-t* state of affairs. If there were post-t* states of affairs that resembled pre-t states of affairs, the former wouldn't be continuations of the latter; they would no more be identical with the latter than the aforementioned new car is identical with the aforementioned old car.

So there wouldn't be *any* reason to identify *any* of U2's occupants with *any* of U1's occupants. In fact, none of U2's occupants would be identical with any of U1's occupants, for the same reason that the previously mentioned new car isn't identical with the previously mentioned vaporized car. Thus, it would only be if U2-states of affairs were continuations of U1-states of affairs that U1 and U2 would be jointly constitutive of some one space-time manifold U, then U2-states of affairs must be continuations of U1-states of affairs.

So nothing in U2 would have anything to do with anything that happened in U1. Nothing in U2 would be an effect or continuation of something that happened in U1. There would be nothing about U2 that connected it to U1. There would thus be no detectable evidence why U2 was a part of the space-time manifold of U1. In fact, there would be no undetectable evidence to that effect. For there would be no undetectable processes---no electron-jumps, nothing too small or too remote to be observable or to have observable consequences---that, if *per impossible* it were known of, would be a reason to hold that U2 was a continuous of some pre-t* universe.

So *operationally speaking*, U2 would not be a continuation of U1. Operationally speaking, U1 would be as distinct from U2 as a world where Dick Cheney is president in 2014 is from our world. Relative to U2, that is to say, U1 would simply be a hypothetical universe; and, should that hypothetical universe exist, the relation borne by it with respect to U2 is the same as the relation borne by a universe in which Cheney is president in 2014 with respect to our universe.

This doesn't mean that U1 doesn't exist. It means that, operationally speaking, it is no more a constituent of any universe of which U2 is a constituent than a 2014 Pres. Cheney Universe is a constituent of a universe of which our universe is also a constituent.

The justification for Inference to the Best Explanation (continued)

This reasoning can be taken further. Suppose that that there are non-initial, U2-internal events $e_1 \ldots e_n$ that occur *ex nihilo*, i.e. that there are events $e_1 \ldots e_n$ that happen in U2 that happen after t*, the interval between them and t* between non-null, that don't have any roots in preceding U2-internal events. ("*Ex nihilo*"—Latin for *from nothing.* An *ex nihilo* event is one that has no roots in anything preceding it, its occurrence be no more favored by the preceding state of the universe than its non-occurrence.) Let t# be the time at which each of $e_1 \ldots e_n$ begins.

Here there are two possibilities:

(1) No U2-internal events that begin at t# have any pre-t# basis.
(2) Some events $e^*_1 \ldots e^*_m$ that happen at t# have a pre-t# basis.

If (1), then the post-t# space-time manifold---let U3 be this manifold---bears the same relation to U2 that U2 bears to U1, in which case, operationally speaking, U2 and U3 are no more jointly constitutive of some one space-time manifold than are U1 and U2.

If (2), then there are events $e^\wedge_1 \ldots e^\wedge_o$ such that $e^\wedge_1 \ldots e^\wedge_o$ precede t# and eventuate, by way of continuous processes, in post-t# events, then U3's relationship to U2 is *not* the same as U2's relationship to U1. Under that circumstance, U2 and U3 *do* jointly constitute a single space-time manifold. But that does not by itself guarantee that $e_1 \ldots e_n$ are constituents of that manifold. Suppose that $e_1 \ldots e_n$ have no effects on any events in U2 or U3. In that case, they neither cause, nor are caused by, U2-internal events or U3-internal events; and in that case, for now familiar reasons, nothing in U2-U3 has anything to do with them and, operationally speaking, $e_1 \ldots e_n$ don't exist as far as U2 and U3 are concerned.

By the same token, if, operationally speaking, $e_1 \ldots e_n$ are occupants of some manifold of which U2 is a constituent, then there must be events $e_{\%1} \ldots e_{\%p}$ such that

(a) $e_{\%1} \ldots e_{\%p}$ begin during U2 and are connected, by continuous event-series, with events in U3

and such that

(b) $e_1 \ldots e_n$ affect $e_{\%1} \ldots e_{\%p}$.

If the only events influenced by e_1-e_n are events that do not themselves have affects on events, such as $e_{\%1} \ldots e_{\%p}$, that are successors of U2-internal-events, then $e_1 \ldots e_n$, as well as its causal successors, are simply divorced from U2 and also from U3. For, under the circumstance just described, the event-class consisting of $e_1 \ldots e_n$ and its causal successors have no effects on anything that interacts with anything beginning in U2 or anything affecting anything beginning in U2; because of which, no event falling into that class is, operationally speaking, part of the same space-time manifold as U2.

Finalizing the justification for Inference to the Best Explanation

Thus, the supposition that there are *ex nihilo* events that are U2-internal is predicated on the assumption that those events are causally connected to U2-internal events that do not occur *ex nihilo.*

This means that, if we know e and e* to belong to the same space-time manifold, then we know there to be some causal connection between them; we know that either

(i) e causes e;*
(ii) e causes e; or*
(iii) some third event e# causes, or is caused by, both.

If observation apprises us of such a connection, then, as far as e and e* are concerned, we don't need to posit unobserved continuities. If observation does not apprise us of such a connection, we do have to posit hidden continuities, since the non-existence of such continuities is incompatible with the datum that e and e* are part of the same space-time manifold.

"But how do you know that they are both part of the same space-time manifold? Whence this 'datum' of which you speak?"

Because we observed each and were therefore affected by each. Since causally connected events are *ipso facto* co-occupants of some one space-time manifold, e and e* are such co-occupants.

Suppose we posit some continuous event-series E such that e causes e* by way of E. If E actually exists, then each segment s of E that lasts for a non-null time-period must satisfy the following condition: s's existence/occurrence cannot be incompatible with that of event e! that is known to occupy the same space-time manifold as e and e*.

If s occupies the same place-time time as e! but is qualitatively different from e!, then E is the wrong series to posit. (It may still be that e causes e*; but if so, e causes e* by way of some event-series that is qualitatively different from E.)

This requirement has significant consequences. e and e* aren't the only events whose causal connection to each other is unknown. And, among the other event-pairs whose causal relation to each other is unknown, there is likely to be one consisting of events e_1 and e_2 such that the event-series linking e_1 to e_2[16] intersects the event-series linking e to e*.[17] This means that, if e$ is the event that connects e_1 to e_2 and constitutes the intersection-point of the series connecting e_1 and e_1 and the series connecting e and e*, then e$ must be continuous with its predecessor and its successor in each of two distinct event-series. By analogous reasoning, for any given segment s$ of series connecting e_1 and e_2, it must be continuous with its predecessor and successor in each of two series.

This means that to validate the datum that two events, e and e*, are part of the same space-time manifold, it is necessary to posit unobserved events that link those two events. There are obviously many event-pairs there are that are known to belong to the same space-time manifold but whose link to each other is not given to observation. In fact, given two non-simultaneous events, in order to validate the datum that they belong to the same space-time manifold, it is necessary to posit a continuous process CP connecting them (either directly or way of some third event to which they are both related as cause/effect) such that every segment of CP is identical

with the segment of each of the infinitely many event-series that intersects with CP.

Herein lies the justification for positing unobserved continuities to explain observed events in terms of other observed events.

The Concept of Natural Law in relation to the problem of induction

Given only that there is a causal connection of some kind between any two co-occupants of some single manifold, how does it follow that the causal connections between past (known) events will bear any resemblance to future (or otherwise unknown) events? How does it follow that the specific IBE's that we make are ever at all principled?[18]

We have seen that the integrity of the space-time manifold depends on the integrity of the causal connections hosted by that manifold. A universe of rogue events, none causative of the other, isn't a universe at all. To the extent that events are causally independent of each other, there is no sense in which they co-occupy some single manifold.

Hume's attempt to delegitimize induction presupposes that the laws of physics can just be suspended. That presupposition in its turn assumes that the laws of physics are separable from the space-time manifold; it assumes, in other words, that they are not *constitutive* of that manifold.

That assumption is erroneous. Suppose that one or more laws of physics were changed. First of all, it wouldn't be possible to change just *one* of these laws. Given a universe in which the speed of light were framework-relative, or in which the Inverse Cube Law were operative, or in which any single one of the laws operative in our universe were replaced by some other law, a number of other laws that are operative in our universe would also have to go.

Laws are not independent of one another. In actuality, gas-pressure is directly proportional to temperature. If it were inversely proportional to temperature, every one of the laws of mechanics would have to be jettisoned or heavily modified.

If the speed of light were framework-relative, then mass and energy would no longer be capable of being converted into each other, at least not in accordance with the same laws that govern such conversions in actuality.

Also, if light did not have a framework-constant velocity, all of our metrical conventions would be thrown off. There would be no principled ways of determining position or, therefore, velocity or, therefore, rates of growth or expansion.

Bearing in mind the interdependencies among natural laws, suppose that at time t* the laws of physics changed. This would have two consequences.

First, relative position, velocity, mass, temperature, and every other physics-relevant parameter, are ascertained by means of procedures that presuppose that certain laws are operative.

Second, the values of those parameters---not our ability to determine those values, but those values themselves---are defined in terms of the laws of physics that are operative.

Thus, if at time t* the laws of physics changed, then the positions, masses, velocities, temperatures, etc. of many, if not all, of the occupants of the post-t* universe (U2) would be discontinuous from their pre-t* (U1) positions, masses, etc. For already stated reasons, the state of the post-t* universe would not be a causal consequence of the pre-t* universe. (Pre-t universe=U1. Post-t* universe=U2.) This means that, if the laws of physics changed at time t*, the result would be the obliteration of the space-time manifold---the universe---in which those laws were operative.

Different laws, different spatiotemporal manifold. Different spatiotemporal manifold, different universe.

The space-time manifold isn't an empty vessel that can house different laws; it is defined by the laws that govern its occupants.

But for the moment, let us suppose *arguendo* that, even though the laws change at t*, the space-time manifold whose laws have been replaced continues to exist. In that case, we have to ask: What were the laws that governed the changes that happened during t*-inclusive intervals? Let's suppose that, at some pre-t* time T, a bowling ball starts rolling. Up until t*, the laws governing its motion are the standard laws of mechanics.

Granting these suppositions, imagine that, during some (short) t*-inclusive interval INT, the ball travels from A to C and that, at t*, it occupies point B.

What are the laws that govern the ball's motion during INT? What are the laws that govern the ball's motion when it is at B? In fact, at t* and after t*, what are the laws that bind together the various ball-constitutive molecules? If so much as a single law of physics has changed at t*, practically every other law of physics will have to change as well. If the Inverse Square Law (ISL) changes, then the mutual attractions of all of the ball's constituent particles change. Of course, the cohesive forces in question are largely non-gravitational; but, owing to the change in ISL, the relative masses of those particles would change, and the cohesive forces in question would have to change if the ball were to continue to exist in anything like its previous form, let alone continue rolling; and it is an open question, certainly, whether they changes would be such as to validate our desire that the ball continue to exist and, what's more, to continue to roll. And, of course, what is true of the ball would be true of the bowling alley and bowling facility and the streets leading to that facility.

Whatever the consequences would be if at a given time even one law of physics were modified, it can be assumed that the consequences would be so drastic that

> *(i) Objects' pre-t* relative positions, masses, etc. would change discontinuously at t*---or, at least, that the likelihood that they would change continuously, or remain constant, at t* would be vanishingly close to zero;*

and

> *(ii) That pre-t* objects would come to an end at t*, as the laws that kept together would no longer be operative.*

We have already seen that, if post-t* events are discontinuous with pre-t* events, then there is no single space-time manifold that U1 and U2 jointly constitute.

We have also seen that in some cases it is a datum that each of two events, e and e*, both co-occupy some one such manifold. This is a datum, as we have seen, because, if a given person is to know of each of two events, each of those events must have affected him; and if each of two events affects a given thing, all three things are *ipso facto* part of some one space-time manifold.

Hence the justification for U (the Uniformity Principle). According to U, the laws of physics don't change. Hume says that any assurance we have as to U's veridicality is induction-based and, therefore, that we have no such assurance. But Hume is not right about that. We know that U is true because, if it weren't, nothing would persist---literally---include the very entities that make inductive inferences. U is a precondition for the making of inductive inferences; so as long as makers of inductive inferences, U holds and inductive inferences either go through or, if they fail, they fail for reasons of incompetence (failure to identify the relevant parameters, failure to compute the relevant sums, etc.), and because of anything having to do with the legitimacy of inductive inference *per se.*

Objects as causal series [19]

Objects (rocks, trees, stars: stable, dynamically unified event-patterns) are systems of interacting parts. The laws governing such systems would cease to exist at t. The same is true of the laws in terms of which the concept of an electron is to be understood. If the laws of physics were different, there would be no electrons, since to be an electron is *inter alia* to be governed by certain laws. The same is therefore true of *every* one of the natural kinds in terms of which pre-t physics understands the world: *proton*, *gravitational attraction*, *distance*, *solidity*, etc.

Objects are constituted by the forces that govern them; those forces are instances of physical laws. Different laws, different forces; different forces, different objects.

What is true of objects is true of all cases of persistence and, therefore, of all instances of causation. Owing to the nomic changes occurring at t, there would be no sense in which anything happening before t could the cause of anything happening after t. Thus, if U1 is the universe before t and U2 is the universe after t, no occupant of U1 would exist in U2. (None of the laws that hold that object together would be operative. None of the laws on whose existence that of its parts depended would be operative after t. No event-series that begin before t would remain in existence after t. (The laws governing that event series would be gone. So would the laws governing its constituents. So would the laws in terms of which mass, velocity, temperature, conductivity, and every other relevant parameter.) It isn't that the series in question would be severed; it's that it would no longer even make sense to suppose that it continued to exist after t: there would no longer exist physical laws relative to which the very existence of that series could be coherently affirmed. This means that U2 would not follow U1. (For U2 to follow U1 would be for some pre-t event to cause some post-t event, which, as we have seen, is not possible.) Thus, given that the laws of physics change at t, then there would be no space-time manifold of which both U1 and U2 were constituents, and the supposition that the laws of physics change at t is therefore tantamount to the supposition that the space-time manifold ceases to exist at t.

This means that *if* there are physical laws, then, until time comes to an end, our inductive practices have a legitimate basis. It also means that, if there are physical laws, then inductive inferences that do not concern the future can have a legitimate basis. Many inductive inferences do not concern the future; they concern what is present or past but is unknown. If I inductively infer that, at this exact time, some object is still solid, that inference doesn't concern the future. If I inductively infer that, several thousand years ago, an asteroid hit a certain part of the Earth's surface, that inference obviously doesn't concern the future. So, provided that physical laws exist, many inductive inferences are unqualifiedly legitimate.

Chapter 7 Determinism, implementation-indeterminism, and the nature of natural laws

“Determinism” defined

Determinism: The doctrine that how the world is at any point in time fixes how it will be at all later times.
Indeterminism: The doctrine that determinism is false. Equivalently, the doctrine that there is some point in time t such that, for some later time t*, how the world is at fails to determine how the world is, in every respect, at t*.

The world is deterministic to the extent that it is law-governed. Supposing that the state of the universe at t determines the state of the universe at later time t*, that is because there are laws of nature $L_1 \ldots L_n$ such that, given the state of the universe at t, there is exactly one way the universe can be at t*.

The world is deterministic to the extent that it is predictable, but it is not predictable to the extent that it is deterministic.

If there are no laws, and thus no deterministic mechanisms, that link the present to the future, then the future is *ipso facto* open, i.e. there are many ways it can be, and predictions cannot be made.

But supposing *arguendo* the world is *deterministic*, it doesn’t follow that it is completely predictable. Setting aside human fallibility, there are a two reasons why determinism doesn’t entail complete predictability.

The first reason why determinism doesn't entail complete predictability: Interference effects

In order to learn about the world, one must interact with it. If I am to learn about another person's mental states, I must ask him questions. But if I am ask him questions, I will change his mental state. If I ask a non-paranoid person "are you paranoid?", he may well, for that very reason, become paranoid. If I want to learn about the bats that live in a certain cave, which is pitch black, I must enter the cave, which will disrupt the bats, and I must use a light, which will further disrupt the bats.

"But aren't these problems capable of being sidestepped?", it will be asked. "Suppose you know when the bats sleep, and you know the layout of the cave, and are thus able to walk around in it without making too much of a ruckus. In that case, surely, you *can* observe the bats without bothering them or, in any case, without disrupting them in a way that prompts them to behave in ways that are not their natural wont? And suppose that I want to learn about a given person's feelings. If I establish an empathic rapport with him, so that he trusts me, I *can* ask him whether he's paranoid without inducing paranoia in him?"

The fact that such precautions must be taken shows that, at any given time, what one can learn about the world is quite circumscribed; that if it's impossible to know X at time t, it may be possible to learn about X at later time t*. But the thesis that the universe is completely predictable requires that, for any time t, it is possible, at t, to predict how, for any later time t*, the universe will be at t*.

Laplace's Thesis

"But Laplace's thesis," it will be objected, "is that a creature of infinite intelligence could predict how things will be."

Prediction requires both intelligence and knowledge, and there is no conceivable being that acquire complete knowledge of the state of the universe at any point in time.

If one is to have a sense-perception of a given thing, there must be an information-bearing causal process linking that thing to oneself. I can hear you talking because the movements of your vocal chords create airwaves in which the structures of those movements are encoded and which then induce structure-similar disturbances in my eardrums. In order to learn about the subatomic world, one must irradiate it, i.e. flood it with light or otherwise bombard it with radiation of some kind or another. But because light is composed of particles (photons) whose masses are comparable to those of the particles that the microphysicist aspires to learn about, the effect of flooding the microphysical world with light is to change it utterly. If you use a flame thrower to learn about a school of rodents, you will succeed *up to a point*: their charred remains will give you tell you a lot about them: but it won't tell you everything. Where rodents are concerned, fortunately, we do not, in our efforts to learn about them, have to flood them with radiation that damages them so. But where microparticles are concerned, we are not so fortunate.

And if there were causal processes that were sufficiently fine-grained that we could learn about some of the microstructures of which the causal processes at our disposal leave us ignorant, it would *still* be impossible to learn about the microworld. The very existence of such fine-grained processes would guarantee that. For how would we learn about *those* processes? By hypothesis, the most fine-grained, least disruptive processes at our disposal would be those very processes. But if we would disrupt such processes if we tried to learn about them by using them.

The reason I can see you without changing you is that, compared to you, photons are very small. But photons are not small compared to photons; and, whatever parts this hypothetical, super-fine-grained causal process consists of, those parts are not small compared to themselves, limiting what could be learned about such parts. If, in order to see you, I had to flood you with objects that were as massive as you are, I couldn't learn much about you. By parity of reasoning, supposing R to be the most undisruptive, fine-grained sort of causal process in existing, there would only be so much one could learn about R by using R.

It being inherent in the nature of sensory knowledge that it involve a causal connection between subject and object, there is no conceivable being that could have complete knowledge of the microstructure of the world. Therefore, nothing could have complete knowledge of the conditions obtaining at t, for any time t, or, consequently, at any later time t*, given that knowledge of the future is necessarily based on the knowledge of the present.

Knowledge of the future=

Knowledge of initial conditions+knowledge of laws+

The knowledge yielded by making the inferences warranted by the first two kinds of knowledge.

Why the existence of the unknown must be granted

George Berkeley, Ernst Mach, C.S. Peirce, A.J. Ayer, all make the following point: the following point: “Supposing that such and such facts are inherently unknowable, then why grant their existence? By supposition, they might as well not exist, even if they do. So why not redefine the term ‘the world’ in such a way that what falls within its scope includes only what can be known after the operations by which knowledge of it is obtained have been carried out?

One fact which falls within the scope of science is that, in order to learn about the world, we must interact with it. Another such fact is that interactions have effects. These two facts entail that what we can learn about the world at time t doesn’t correspond with perfect precision to the world at time. Given that two facts that do fall within the scope of science entail that there are things about the world that cannot be known, it follows that, for science no less than for metaphysics, the world is to some extent unknowable.

Also, we can know that situations of a given *kind* obtain without there being any one such situation of whose existence we are aware. There has never been a single photon of which any person was aware; but we know that photons exist, since, if they didn't, the photoelectric effect could not be explained.

Second reason why determinism doesn't entail complete predictability: No way to predict how knowledge will grow

No creature can possibly predict what it will come to know. Suppose you know now, at time t, that, during the next 20 years, you will learn truths $t_1 \ldots t_n$, and no other truths. In that case you *now* know $t_1 \ldots t_n$ and they therefore *aren't* what you will learn in the next 20 years. If we replace $t_1 \ldots t_n$ with any other list of truths, the result is the same. Therefore, it is possible to know now what you will know later. Given that what you know affects what you do, and thus affects how the world is, it follows that no creature can know exactly what it will do or, therefore, exactly how what it will do will affect the world.

Supervenience

If x and y are identical in respect of their microphysical properties, then it is inconceivable that they should differ in respect of their biological properties.[20] Wherefore, the biological *supervenes* on the microphysical. In general, one class of facts supervenes on another class of facts exactly if there can't be a change in the former without there being a change in the latter but not *vice versa.*

Supervenience≠identity

Given two pieces of clay that are identical in respect of their strictly physical properties (e.g. their shapes, masses, etc.), it is inconceivable that they should differ in respect of their aesthetic properties. But aesthetic properties cannot be identified with physical (physics-related) properties.

First of all, two objects that are physically different can be aesthetically identical. Given only that x and y differ in their chemical composition, their masses, their degrees of conductivity, etc., it doesn't follow that they differ in their aesthetic properties. The aesthetic is thus *multiply realizable* with respect to the physical. The same is true, for the same reasons *mutatis mutandis*, of the biological.

The biological is thus variable with respect to the physical. A biological mechanism of a given kind can be *implemented* by entities that are physically very different from one another.

Software is variable with respect to hardware. Software-identical systems may be hardware-distinct. (But hardware-identical systems cannot be software distinct.) Software-realities therefore *supervene* on hardware realities.

Psychological supervenience difference from other kinds of supervenience

It is said that the mental supervenes on the physiological: given two creatures that are physiologically identical, it is impossible that they should differ psychologically.

Supposing this true, the way in which it is true is very different from the way in which it's true that biological supervenes on the physical. It is incoherent to suppose that two entities that are microphysically identical should differ biologically.

By contrast, it is not incoherent, though it is counternomic, to suppose that two physiologically identical creatures should differ in respect of their psychological properties.

Bound by physical law≠
explained by physical law

Nothing can violate the laws of physics. There are no mass-energy displacements that are not governed by the laws of physics. For this reason, many have thought that everything (everything spatiotemporal, that is) can be *explained* in terms of the laws of physics.

Spurious reasoning for a false conclusion. The statue is not identical with the clay of which it consists. They have different modal properties (one can destroy the statue without destroying the clay); so the clay isn't identical with the statue. (The clay *constitutes* the statue; constitution isn't identity.)

By parity of reasoning, computers aren't identical with heaps of atoms. The laws of physics tell you how those atoms behave. They don't tell you anything about the computer. The laws of physics don't say anything about computation, storage space, etc. Those concepts (*computation*, etc.) don't exist as far as those laws are concerned. So a physics-omniscient being does not in virtue of that fact have any computer-specific concepts and therefore does not have any computer-related knowledge.

Two software-identical units may be physically very different. So given only that x and y are software-identical, nothing can be said about their respective physical realizations. Consequently, there is no way to map software-truths onto hardware-truths.

'But can't it be said that x and y will resemble each in certain abstract, functional respects?'

Yes, but physics is blind to such similarities. The similarities in question are to be understood in terms of concepts (e.g. *computation, recursion, algorithm*) for which the laws of physics make no provisions.[21]

Software-determinism≠
statistical determinism

Suppose that a given computer x has a deterministic program. (‘x has a deterministic program’=‘there exists a set S of inputs such x, when not malfunctioning, assigns no more than output to any given member of S.’) Computers break down, and there is thus ‘machine error.’ Should this happen to x---should x sometimes fail to say ‘4’ when asked ‘what is twice two?’---the conclusion to draw is that x isn’t implementing its program, not that its program is indeterministic. To the extent that x is a computer---to the extent that x is something that implements a program---x is a deterministic system. It is only to the extent that x is failing to implement a program that x is indeterministic. So it’s only to the extent that x isn’t a computer that x is indeterministic.

Implementation-determinism≠
program-determinism

Where higher-order laws are concerned, determinism≠statistical determinism.

Given only that biologically identical inputs may be followed by biologically distinct outputs, it doesn't follow that the laws (if any there be) of biology are 'indeterministic' or 'probabilistic.' What follows is *either* that those laws are indeterministic *or* that organisms imperfectly operationalize their biological programs. Given only that economically identical situations may lead to economically distinct situations, it doesn't follow that the laws of economics are statistical. And so on.

Correlatively, lower-order indeterminism≠higher-order indeterminism. Supposing that the laws of physics are indeterministic, it doesn't follow that higher-order laws are indeterministic. Supposing that $x_1 \ldots x_n$ are indeterministic with respect to the laws of physics, it doesn't follow that they are indeterministic with respect to this or that biological/computational/etc. program.

Supposing *arguendo* that the laws of microphysics are indeterministic, it doesn't follow that *aggregates* of microparticles behave indeterministically, for the reason that the various norm-non-compliant behaviors on the part of individual particles cancel one another out.

Probability in relation to higher-order mechanisms

Let W1 and W2 be two universes satisfying the following conditions. In each, the very same laws are operative. In each, the microphysical laws are indeterministic. In each, the fact that those laws are indeterministic has the consequence that it is very unlikely that computers (or cars or any other large-scale device) will always do what it is supposed to do: given those laws, it is unlikely, in either W1 or W2, that every computer will always comply with its program and will never, as a result of machine-error, produce an input-inappropriate output. But in W2, by sheer coincidence, every computer always does what it is supposed to do. Moreover, in W2, every computer does what it is supposed to do *even though*, inevitably, every computer eventually breaks down: those computers that break down simply happen, by sheer coincidence, to produce program-compliant results, even though, their hardware-conditions being what they are, it is unlikely that they should do so.

In W2, the likelihood that a given computer will behave in a program-compliant doesn't exceed the likelihood that a given computer in W1 will do so: the two probabilities coincide. But in W2 the previously-described rash of coincidences obscures the deterministic, and therefore that probabilistic, structure of events. In W2, it *seems*, in consequence of those coincidences, that computers are 100% likely to behave in a program-compliant fashion 100% of the time, even though, in fact, it is not 100% likely---even though in fact their *always* behaving in a program-compliant fashion is almost 100% *unlikely* (granting that their *usually* behaving in a program-compliant fashion is 100% *likely*).

Statistical trends reflective of, but distinct from, probabilities

The Frequency Theory of Probability (FT) tells us that, in W2, the probability of a computer's never malfunctioning is 100%; in any case, the FT tells us that the W1-likelihood of computer-error exceeds the W2-likelihood.

This is plainly false. In each universe, the likelihood of computer-error is the same, even though, in W1 the likelihood *seems* high and, in W2, the likelihood *seems* low and even though, according to FT, the first likelihood *in fact* exceeds the second. What is going on is that, in W1, the statistics are consistent with the deterministic-probabilistic structure of the operative laws, whereas in W2 they are not; so that, in W1, the statistics lead one to the right result, whereas the statistics fail to do so in W2.

Evidence of probabilities vs. actual probabilities

Statistics are *evidence* of probabilities; statistics are not *themselves* probabilities. According to FT, statistics are identical with probabilities; therefore, FT is false. Also, statistics are only as truth-conducive as they are consistent with the causal mechanisms that generate them. If, by coincidence, the statistics should be out of line with the underlying causal realities---if, by sheer coincidence, not one person who smokes gets cancer or emphysema---then the statistics simply aren't good evidence of the relevant mechanisms or, therefore, of the relevant probabilities.

In a world governed by the same biological---and non-biological—laws as ours, but in which by sheer coincidence smokers never got cancer/emphysema, doctors would be no less wrong than they would in our world to tell their patients that smoking was harmless. In such a world, it would be understandable, given the available evidence, that doctors should take the view that smoking was harmless: but that view would be no less wrong than it is here, and the probability, in that world, that a given smoker would succumb to cancer/emphysema would coincide with the corresponding probability in our world.

Assume that M is brand-new and that it is well-constructed. Bearing this in mind, consider the statement that:

(L) The likelihood that M will respond to a given input with an appropriate is high, but not 100%.

Let W be a world in which, by coincidence, computers always behave in a program-compliant manner but which, within the limits set by that qualification, is otherwise just like our universe. In our world, the likelihood that a brand-new PC will behave in a program-compliant manner is close to, but short of, 100%.

If FT is right, then the probability in W of a computer-malfunction is 0%, even though the actual probability is obviously the same as the corresponding probability in our universe and is therefore greater than 0%.

Conclusion of Chapter 7

When it is said that that probability of a computer-malfunction is greater than 0%, the statement being made is to be interpreted causally, not statistically. What is being said is that, the laws of physics being what they are and the construction of the computer being what it is, it is a possibility, albeit one that may never be actualized, that brand-new/well-constructed/etc. computers malfunction.

Chapter 8 Measurement

Measurement vs. enumeration

Given any two lengths, there are intermediate lengths. There are therefore no minimal units of length. Consequently, given an object L, there is no class C of objects such that by counting C's members, one thereby determines L's length.

Even if L is known to be length identical with some object L* that consists of indivisible particles $p_1 \ldots p_n$ arranged in a single file, it must be known that each such particle is length-identical with each other such particle if it can be legitimately said that either L* or L is n-units long. And given any such particle, there is no class C* such that the length of that particle can be determined by counting C*'s members, since, by hypothesis, each such particle is non-composite.

Also, even if each such particle is length-identical with each other such particle, it must be known that length is *strictly additive* if it is to be known that the length of a $p_1 \ldots p_n$ is n-units long. In other words, it must be known that, if a straight object consists of one segment that is x units long and another that is y units long, that object is x+y units long.

As it turns out, length is *not* additive. Length is approximately additive, when the lengths in question are small. It is not approximately additive when they are big.

The very fact that length is non-additive shows that, even if it were additive, "L is n units long" is not equivalent with "L consists of n unit-length segments": those statements have different meanings.

Also, supposing that length were additive, no act of enumeration would disclose that fact: enumeration results in principled length-determinations only if it's *already* known that length is additive.

Relativization to a standard the essence of measurement

When it is said that L is longer than L*, L's relative length is being given---its length relative to L*.

When it is said that L is two meters long, it is *still* L's relative length that is being given---L's length relative to some *standard.*

1. "Joe is as tall as Bill"

tells you nothing about Joe's height, or Bill's height, unless you happen to have context-specific information.

2. "Joe is 6 ft tall"

tells you how tall Joe is, even if you lack context-specific information.

Nonetheless, (1) is a comparative statement, just (1) and just like:

(3) "Joe is taller than Mary."

The difference between (2), on the one hand, and (1) and (3), on the other, is that (2) relativizes Joe's height to a universal standard, whereas the others relativize it to a context-specific standard.

The essence of measurement is comparison with respect to a standard. This is true of every kind of measurement.

Relative vs. absolute length

There is no way to determine an object's length *except in terms of the length of some other object.* When an absolute length is assigned to some object L, it is being said that there is some object LS such that

(a) LS is stipulated to have a length of one unit

and such that

(b) L's length compared to LS is n units.

Let LS be some object---e.g. a platinum bar in a temperature-controlled room---that is stipulated to be one meter long. If, when juxtaposed with L, in the way previously described, each of L's endpoints is contiguous with each of LS's endpoints, then L and LS are the same length. Otherwise, one of them is longer than the other.

Let's suppose that L is longer than LS. In that case, there are two possible scenarios:

(1) There are whole numbers p and q such that the ratio of L's length to LS's length is p/q, i.e. the ratio of L's length to LS's length is given by a rational number.
(2) There are no such numbers, i.e. the ratio of L's length with respect to LS's length is not given by a rational number.

In case (1), there is some number n such that, if L is divided into n segments of equal length, L* being an arbitrary one of those segments, LS's exact length can be established by the method earlier described. That length will be an improper fraction n/m, such that L is m times the length of L*.

In case (2), if L is divided into n length-identical segments, L* being an arbitrary such segment, there is some integer m such that, the larger n is, the smaller the difference between LS's exact length, on the one hand, and the measurement yielded by the method just-described.

Three axioms of measurement-theory

(i) If L_1 is longer than L_2, and L_2 is longer than L_3, then L_1 is longer than L_3.
(ii) If L_1 and L_2 are length-different, one is longer than the other.
(iii) If L_1 is the same length as L_2, and L_2 is the same length as L_3, then L_1 is the same length as L_3.

No viable system of length-determination is inconsistent with (i)-(iii).

Given any measurable magnitude m, what is true of length-determination is true of m-determination. In (i)-(iii), replace the references to length with references to temperature or mass or any other measurable magnitude: the resulting statement-set is true.

Measuring an object's length involves assigning a real number to its length: same length, same number; different lengths, different numbers. If any one of (i)-(iii) is violated, length-identical objects will be assigned different numbers and length-distinct objects will be assigned the same number. Number-assignments are therefore meaningless unless they are consistent with (i)-(iii).

Scalar quantities vs. vector quantities

Length is a *scalar* quantity. A scalar quantity is one that is given by a single number, as opposed to an ordered n-tuple.

A vector quantity is one that is given by an ordered pair or, more generally, an ordered n-tuple. Velocity is a vector quantity; an object's velocity is given an ordered pair of numbers, one corresponding to speed, the other to direction.

In this context, we will focus only on scalar quantities. The principles underlying the determinations of vectors are obvious extensions of those involved determinations of scalar quantities.

Mass, temperature, density, viscosity, pressure, speed, conductivity, and malleability are all scalar quantities. Speed=distance/time. But speed is a scalar quantity. (Quotients are single numbers, not ordered n-tuples.) Density=mass/volume. But an object's density is a scalar quantity.

Scalar magnitudes identical with one-dimensional manifolds

Even though a piece of paper is a four-dimensional object, it is a one-dimensional object *qua* thing that has a length. Any given length-determinate entity is, to that extent, a one-dimensional entity, even though any given occupant of space-time is four-dimensional.

Given a scalar magnitude m and any object x, *qua* m-determinate entity, x is one-dimensional, even though x *per se* cannot possibly be one-dimensional.

Congruence

If length-determinations are to be made, it must be said what it is for two objects to have the same length. Two objects have the same length if they are length-congruent. Two objects are length-congruent if (but, given that there are many ways to establish length-congruence, not *only* if) two conditions are met, namely:

(i) When juxtaposed, each endpoint of the one is contiguous one of the endpoints of the other,

and

(ii) The distance between those two objects is constant---it doesn't wax and wane.

L is longer than L* if L* is congruent with a proper segment of L.

Absent a definition of congruence, "L is longer than L*" and "L and L* are length-identical" are devoid of meaning. It isn't that their truth-values cannot be known. It's that there are no truth-values to know.

Given any magnitude m, "x and y are m-identical" and "x m-exceeds y" are devoid of meaning. But for now, we will focus on length.

Absent a definition of length-congruence, no object L can be a length-standard. For L to be a length-standard, it must be said what it is for L* to be length-identical with L.

But *given* a definition of congruence, there is no object L such that "L* is one unit long" means "L* is length-identical with L" and thus no object L such that "L* is n units long" means "L* is n-times longer than L."

Explanation: Let L be some specific meter-rod that is taken as a standard. Given some L-congruent object L*, L* is capable of serving as a length-standard. It may even be a better length standard than L. It may turn out that, if L* is assumed to be length-constant, conditions (i)-(iii) are satisfied, but that those same conditions are violated if L is assumed to be length-constant.

n-dimensional congruence

Length-congruence holds between one-dimensional objects; triangle-congruence holds between two-dimensional objects; pyramid-congruence holds between three-dimensional objects. For any integer n, there are (hypothetical, if not actual) n-dimensional objects and, consequently, n-dimensional congruence.

The use of provisional metrical standards and approximate physical laws

A given object is provisionally taken as a length-standard. A congruence-relation is defined. On that basis, other objects are measured. Those measurements yield knowledge of physical laws. Those physical laws entail that, if an object of a certain type is taken as a standard, the result is a consistent set of measurements and also that an inconsistent set of measurements results if an object of some other type is taken as length-standard.

Given any magnitude m, if the references to length in the statements just made are replaced with references to m, the resulting statements are true.

Triangulation and precisification

Small distances can be measured using meter rods, and they can also be measured using light rays. Large distances can be measured using light rays, but they cannot be measured using meter-rods. Within narrow horizons, meter-rod measurements and optical measurements yield the same results.

This is a contingent fact. Suppose it turned out---as it well might have---that the numbers yielded by meter-rod measurements could not be coordinated with the numbers yielded by optical measurements. In that case, there would be no one magnitude that was measured in both ways, and statements to the effect that the Sun is n meters from the Earth would be meaningless. But such statements are not meaningless, since, within the horizons within which both methods of measurement can be used, the results generated by the one method correspond to those yielded by the other.

Precisification in relation to magnitudes other than length

One can use a thermometer to measure the temperature of bath water. One cannot use a thermometer to measure the temperature of magma, let alone the temperature of a star.

On the basis of thermometer-based measurements, one acquires knowledge of low-level physical laws (e.g. that such and such metals have such and such coefficients of expansion). On the basis of *this* knowledge, one acquires knowledge of higher-level physical laws (e.g. that radiation-intensity varies inversely with the square of the distance from the heat-source). On the basis of the totality of this information, one can compute very high temperatures.

We know from personal experience that

(i) A stronger flame is needed to boil a given amount of water in a short period of time than over a longer period of time;
(ii) It is easier to hoist a given container when it's empty than when it's full; and
(iii) Cars and computers that aren't maintained become dysfunctional more quickly than those that are maintained.

On the basis of these facts, we know that

(i) An object's mass correlates with the amount of energy needed to heat it up;*
(ii) An object's density correlates with the amount of energy needed to displace it; and*
(iii) More energy is required to sustain a low-probability situation (e.g. an assemblage of objects that constitutes a working machine) than a high-probability situation (e.g. an assemblage of objects that serves no purpose).*

On the basis of (i*)-(iii*), we can determine the sizes, masses, distances (from Earth and from each other), and the ages of celestial bodies. And on the basis of *this* knowledge, we can identify precise laws, to replace the approximate ones with which we began.

Summary of findings

Day-to-day experience gives us approximate metrical knowledge, on the basis of which we can identify correlations on the basis of which we can obtain hitherto unobtainable metrical data. On the basis of this metrical data, we can identify precise laws, on the basis of which we can both add to and refine such metrical knowledge as we already have.

Synchronic vs. diachronic measurements

Two meter-rods can be laid side by side. Non-simultaneous events cannot be laid side by side. If it is to be known how two such events compare in respect of duration, one must know of some periodic process (some event-sequence, each of whose members is of the same type as its successor: e.g. the round-trips made by a beam of light between two mirrors, rotations of the earth or of a cesium atom, somebody's heartbeat) that occurred during each event and, further, one must know how many installments of that processes occurred during each. If m installments of that process occur during the one event and n installments occur during the other, then, if that process is indeed a perfectly regular one, the length of the one relative to the other is n/m.

Metrical Conventionalism

"But it cannot be said," it will be objected, "that a given periodic process is regular unless there is already a way of determining duration. And there is no way of determining duration except in relation to some periodic process. Therefore, some periodic process must be *stipulated* to be perfectly regularity: it must be perfectly regular *by convention.*"

This position, duly generalized, is known as *conventionalism.*

Conventionalism is false. This will be discussed below. Long story short: If duration is measured in terms of a process that is intuitively known to be erratic (e.g. somebody's heartbeat[22]), the resulting measurements aren't consistent with one another, but duration-measurements *are* consistent with one another if duration is measured in terms of a process that is intuitively known to be regular (e.g. the series of roundtrips made by a light beam between two mirrors whose distance from each other is fixed). Given that, relative to a given framework, an event cannot have distinct durations, it follows that conventionalism is false.

Extensive vs. intensive magnitudes

If 90° water is added to 70°water, the result is not 160° water. It is water that is n°, $70 \leq n \leq 160$.

If a one lb. object is put on top of a two lb. object, the result is an object whose weight exceeds the weight of either of the two original objects. (The resulting object weighs approximately three lbs.)

The length of a straight object is (approximately) the sum of the lengths of its segments.

Length and mass are *extensive* magnitudes. Temperature is an *intensive* magnitude. "Intensive" is synonymous with "non-extensive." An extensive magnitude is one such that the degree to which a given object has that magnitude exceeds the degree to any given one of its parts has that magnitude. An *additive* magnitude is one such that the degree to which a given object has that magnitude is the arithmetical sum of the extent to which each of its parts has that magnitude.

Length and mass are extensive, and within limits they are both additive. Speed is extensive and, within limits, additive. Intelligence is intensive. So are temperature and density.

If two three ft. objects are laid side by side, the resulting object is not six ft. When it is said that length is extensive, what is meant is that there exists a non-contrived, knowledge-conducive way of conjoining objects such that the resulting object's length exceeds any given one of the objects conjoined. And when it is said that length is additive, what is meant is that there is exists a non-contrived, knowledge-conducive way of conjoining objects in such a way that the resulting object's length is the arithmetical sum of the lengths of the objects conjoined.[23]

Relativity Theory and the non-additivity of scalar quantities

If you are at a train-station, your speed relative to that of a light-beam is 0 mph. If you get on a train traveling in the direction of that light-beam, your speed relative to it is still 0 mph. If you get on spaceship that is traveling in that direction at a rate of 150,000 miles/second, your speed relative to it is *still* 0 mph.

Given any framework F, the speed of light relative to F is c (≈186,000 miles/second). A consequence---actually, the very essence of Special Relativity---is that there is no optical test by means of which one determine whether or not one is moving.

If speed were additive, these statements would be counterlogical. But speed, though extensive, is not additive. If you are moving at a rate of 30 mph on a train that is moving in the same direction at a rate of 30 mph, your speed relative to the train station is less than 60 mph.

A corollary is that length is non-additive. A corollary of this corollary is that volume is not additive.

Direct measurements always yield rational numbers

Suppose that the ratio of L's length to LS's length is √2. In that case, there is *ipso facto* no number n such that, if L is divided into n-equal parts, LS's length is some integral multiple of n. So if L's length, relative to LS, is √2, that fact cannot be established by measuring L. If a given object's length is given by an irrational number, that object's length cannot be established by measuring that object.

Nonetheless, it can be known that a given object's length is given by such a number. If there is a triangle T such that

(i) L is T's hypotenuse,
(ii) LS is one of the other sides, and
(iii) The third side is length-identical with LS, then L's length relative to L is √2.

If there is a circle C such that L is length identical with C's circumference and LS is length-coincident with C's diameter, then L's length, relative to LS, is π.

Nomic Conventionalism

Suppose, as before, that LS is our length-standard, i.e. that it is stipulated to be length-constant, its length being one unit. It obviously isn't by comparing LS with some other object that we know LS to be length-constant. According to many (e.g. Russell, Poincaré, Einstein), this entails that, in taking LS to be length-constant, we are simply adopting a convention, not registering a fact. Their argument:

> (CM) Since LS is our length-standard, there isn't some other object Ln such that LS's having a given length, or its being length-constant, was established by comparing LS with Ln. Given that length is comparative length, there is no other way to establish that LS is length-constant or that it has a given length. Therefore, in taking LS as a length-standard, we are simply adopting a convention: we are not acknowledging a fact about LS's length. Therefore, the statement "LS is one unit long" merely records a convention; it doesn't express a fact. The same is therefore true of "LS is length-constant (i.e. there is some length l----in this case one unit----such that, at any given time LS's length is l).

For argument's sake, suppose that CM is cogent and that it cannot possibly be on the basis of known facts about LS's length that we take LS to be our length-standard. In that case, given *any* object L!, L! has just as much right as LS to be taken as a length-standard. L! can be an eel, a cylinder-shaped balloon, a sausage in a frying pan---anything at all. So far as there is any reason to take LS as a length-standard, and thus as length-constant, those advantages are of a purely practical nature. It may be that, by taking LS as length-constant and regard (cylinder-shaped balloon) L! as length-variable, the resulting system of physics is relatively simple; and it may be that if we take L! as length-constant and thus regard LS as length-variable, the resulting system of physics is prohibitively complex. If so, it doesn't follow that the first system of physics is more correct than the second. If $L_1 \ldots L_n$ are the laws constitutive of the LS-based system of physics and $L^*_1 \ldots L^*_m$ are the laws constitutive of the L!-based system, $L^*_1 \ldots L^*_m$ are no less real than $L_1 \ldots L_n$. The question whether a given physical law exists has no meaning except in relation to objects that are taken as standards. Since it is entirely a matter of convention whether a given object is such a standard, the laws of nature are determined by our conventions.

This doctrine, though distinct from the one same-named doctrine previously considered, is known as *conventionalism.*

Thus, "conventionalism" is ambiguous. Sometimes it refers to *metrical* conventionalism, and sometimes it refers to *nomic* conventionalism.

Nomic conventionalism is seldom distinguished from nomic conventionalism. But those who do distinguish between them almost always take nomic conventionalism to be a consequence of metrical conventionalism. Their argument:

> If nomic conventionalism is correct, it is only relative to some arbitrary choice of standards that the so-called laws of physics (the Boyle-Charles Law, the Inverse-Square Law, the Parallelogram Law, $e=mc^2$, and so on) are accurate. If balloon L! is taken as a length-standard, none of these so-called laws is even a truth, let alone a truth that expresses a law. These so-called laws *are* laws if optical standards of

length (and length-related magnitudes: e.g. speed, distance, temperature) are adopted.

This argument refuted

Supposing *arguendo* that conventionalism is correct, it doesn't follow that the laws of nature are *entirely* a matter of convention. What follows is that laws of nature are *relations* between convention-based standards, on the one hand, and empirical data, on the other.

Given only that LS is taken as a length-standard, nothing follows as to what the laws of nature are. But given that LS is taken as a length-standard, then *supposing that LS-based measurements have yielded such and such results*, it follows that, so far as those measurements were conducive of knowledge of natural law, there are thus and such laws of nature.

So supposing that conventionalism is correct, what follows is that the laws of nature describe non-conventional *relations* between metrical standards, on the one hand, and objects/events, on the other.

Conventionalism vs. Strict Conventionalism

Let *strict conventionalism* be the position that the laws of nature are *entirely* a matter of convention. Strict conventionalism isn't a form of conventionalism. In fact, the two doctrines are incompatible. If the laws of nature were convention-*determined*, as opposed to merely convention-*based*, there would be no need to take measurements to identify those laws.

By the same token, so far as measurements are needed to identify them, they are, at most, convention-based, not convention-determined. The results of measurements are not determined by our metrical conventions. Given only that LS is stipulated to be one meter long, nothing follows as to how long a given car is: that car must be measured, and the outcome of that measurement is an empirical matter.

Supposing that strict conventionalism is false but that conventionalism is correct, magnitudes (e.g. lengths and distances) are properties non-conventional *relations* between events, on the one hand, and metrical conventions, on the other. In other words, if conventionalism is correct, then, even though the statement "the speed of light is framework invariant" doesn't describe a fact about light *per se*, it *does* describe a fact about the relationship between the behavior of light, on the one hand, and our metrical standards, on the other.

Given the fact that, if conventionalism is right, laws hold in virtue of non-conventional relations between conventions and events, it follows, for reasons now to be described, that conventionalism is a self-defeating doctrine.

Conventionalism inconsistent with the desideratum that causal anomalies be eliminated

For argument's sake, let's suppose (probably correctly) that eels periodically increase and decrease in length, and let's suppose that we decide that we take some eel L! as our length-standard. Obviously L! *per se* can only be used in an extremely limited number of contexts. In order to use L! as a standard, there must be objects $L!_1 \ldots L!_n$ such that, for each i ($1 \leq i \leq n$), $L!_1$ is congruent with $L!_n$. Each series member must be D-congruent with each other such member. Further, there must be multiple such series: supposing that $L!_1 \ldots L!_n$ mediates between L!, on the one hand, and some car x in Arizona, on the other, there will inevitably be some object y in Tibet (or on Mars) such that for no i ($1 \leq i \leq n$) is $L!_i$ ever in a position to be juxtaposed with y. And each member of each such series must be D-congruent with each member of each other such series.

The empirical data being what it is, the likelihood is vanishingly close to zero that the length of a given object relative to a given such series-members will coincide with its length relative to every other such member or even any one such member.

A system of measurement must yield consistent results: if, relative to a given metrical standard, a given object (at a given time, relative to a given framework) both has, and does not have, a given length n, then that standard cannot coherently be used. For use of such a standard constitutes acceptance of the absurdity that $n \neq n$, for some number n.

It is an empirical fact that, when objects that we intuitively regard as length-variable (e.g. eels, cylinder-shaped balloons) that are used as length-standards, the resulting measurements are inconsistent with one another.

It is also an empirical fact that, when objects that we intuitively regard as length-constant are used as length-standards (e.g. platinum rods in temperature controlled rooms), the resulting measurements are consistent with each other to a high degree of approximation.

Finally, it is an empirical fact that, if an optical standard is used, there are fewer inconsistencies in the resulting measurement-set than if a non-optical standard is used.

Given that $n=n$, for any n, and given that use of an eel standard constitutes a rejection of this truism, whereas use of an optical standard constitutes acceptance of it, it must be said that, if eel-based measurements yield one set of physical laws and optics-based measurement yields a different set of laws, the latter are the laws that govern nature and the former are not.

Conclusion of the present work

Causal series are continuous series. Information-bearing series are continuous structure-preserving series. To know explanations is to know causes. Scientific hypotheses are correct when they eliminate discontinuities, and they eliminate discontinuities when they themselves are constitutive of information-transmissive continuities.

---

[1] Sometimes the word 'primitive' is used instead of 'direct.' And, as previously stated, the words 'direct' and 'indirect' are, in this context, synonymous with, respectively, 'non-inferential' and 'inferential.'

[2] This example is due to Russell. Cf. his (1914) book *Some problems of Philosophy.*

[3] As J.L. Mackie establishes in *The Cement of the Universe.*

[4] *The Analysis of Matter* (1928) and *Human Knowledge: Its scope and limits* (1948).

[5] This isn't really a law. But it is a principled, as opposed to an accidental concomitance, and it is therefore *law-like* and, for our present purposes, enough like a law that it may in this context be a considered a law.

[6] This isn't really a law. But it's almost a law, and will serve as an adequate law-proxy in this context.

[7] Or, to be precise, only *to the degree* that C caused E, *C's* degree of relevance being inversely proportional to the number of cofactors.

[8] Galileo's Parallelogram: Suppose that x is moving in a rectilinear fashion. First, x is struck by y, causing x to change velocity (direction and/or speed); then x is struck again by z, causing another such change in x. If the forces exerted by y and z had been exerted at once, x's final state of motion---its state of motion after being struck, first by y and then by z---is identical with what its state of motion would have been had it been struck simultaneously by y and z (assuming that, in both cases, the force with each of y and z strikes x is the same). Thus, x's vector is a function of the two forces operating on it, and it is by vector-adding those forces that said vector is determined.

[9] Which, incidentally, is not the case, since there are people who have a disorder whereby they cannot feel pain. Such people, unfortunately, do not live long, since they constantly suffer injuries without knowing it and therefore without doing anything about it.

[10] "Force of inertia" is not a term that physicists ever use. But in this context it is convenient to

use it.

[11] The term “degenerating research program” is due to Imre Lakatos, a philosopher of real intelligence who did meaningful work in the philosophy of science.

[12] Discovered by Carl Hempel; stated in his paper “Paradoxes of Confirmation,” which was republished in *Aspects of scientific Explanation.*

[13] Note: X may be an object, e.g. a particular emerald; an object-type, e.g. the class of emeralds; an event, e.g. a bolt of lightning; an event-type, e.g. the class of bolts of lightning; a structure, e.g. a building o f a particular configuration; or a structure-type, e.g. the class of buildings of a particular configuration.

[14] *Note about the term “Inference to the best explanation"*: Some IBE’s are not the *best* explanation available. Some IBE’s are decidedly poor explanations. But any given IBE is believed by its proponents to be the best available explanation. So “IBE” is not an honorific term: in describing an inference as an “IBE,” nothing is being said as to whether it does in fact eliminate the anomalies generated by some data-set; what is being said only that, in somebody’s mind, it does so.

[15] See Freud (1914), *Totem and Taboo*. Freud’s position: Infants and otherwise undeveloped people project the causal order of their thoughts onto reality, and they consequently believe that the causal dependencies that hold among their mental representations of things hold among those things themselves. For this reason, they believe that, if they think it, it will happen: hence the belief in the ‘omnipotence of thought’ (Freud’s expression).

Freud’s position is not that people are wrong to ascribe *a* causal structure to reality, but that they are wrong to ascribe the causal structure of their own minds to reality. Freud believes that there are laws of physics, etc. But he also believes that the laws of psychology are not the laws of physics and, further, that people---all people, at some level of another---have yet to grasp this fact.

Hume also believes that people erroneously project the psychological-causal order onto external reality. But Hume believes, in addition, that the external world is causally unstructured: there is *no* external causal order. So for Hume, the error is that people are attributing *any* causal structure to the external world, whereas, for Freud, it is that people are attributing a causal order to external world that is different from the one that it actually has.

Hume’s position is incoherent. If Hume believes in psychological causation---that is, if he believes that, in the mental sphere, causation≠regularity---then Hume is *ipso facto* wrong to hold that *in general* causation=regularity, and his belief that, in the external world, causation=regularity is discrepant with his belief that, in other contexts, causation≠regularity.

[16] Either as cause and effect or as common effects/causes of some one effect.

[17] Even if this isn’t true of the series linking e to e*, it is true of other such series. Plus, it is demonstrably incoherent, by (not so obvious) extensions of the points already made, to suppose

it not true of that particular series; and we therefore might as well suppose it true of that particular series.

[18] Sometimes they turn out to be correct in hindsight. But how do we know that there are any cases of IBE's such that, when we were putting forth those IBE's, we had good reason for doing so? According to Karl Popper: We don't. We never have any good reason to put forth this as opposed to explanation. All we can do is put forth *an* explanation and then hope it doesn't turn out to be false.

[19] Russell (1928, 1948) deserves credit for making and justifying this point, using this exact phrase.

[20] If neither has any biological properties---if, for example, they are both rocks—then they don't *differ* in their biological properties.

[21] 'But don't physicists use computers? Don't' they perform computations?' Yes. But the methods used by a science do not necessarily fall within the scope of that science. Botanists study plants; to do so, they use magnifying glasses; but they don't study magnifying glasses. Magnifying glasses fall within the bailiwick of optics, not botany.

[22] This example is borrowed from Hempel (1953).

[23] Additivity (as Russell points out in *The Analysis of Matter*) is not really a property of physical magnitudes, but of cardinal numbers.